Holding Period Effects in Dividend Strip Returns

Benjamin Golez Jens Jackwerth *

November 2023

Abstract

We estimate short-term dividend strip prices from 27 years of S&P 500 index options data (1996-2022). We use option-implied interest rates when estimating strip prices and longer holding period returns to mitigate measurement error. We find that Sharpe ratios for short-term strips are similar to or higher than Sharpe ratios for the market. Short-term strips also have a low market beta and a positive alpha. Over the business cycle, realized term premia (i.e., the difference between market and strip returns) and the term structure of Sharpe ratios move countercyclically, whereas the term structure of alphas moves procyclically.

JEL Classification: G12, G13, G35

Keywords: dividend term structure, dividend strips, option pricing, option-implied interest rates

---

*Benjamin Golez is from the University of Notre Dame, 256 MCOB, Notre Dame, IN 46556, USA, Tel.: +1-574-387-9597, Email: bgolez@nd.edu. Jens Jackwerth is from the University of Konstanz, PO Box 134, 78457 Konstanz, Germany, Tel.: +49-(0)7531-88-2196, Email: jens.jackwerth@uni-konstanz.de. We thank Ralph Koijen (editor), two anonymous referees, Tobias Sichert (discussant), and participants at the ESADE Spring Finance Workshop 2023.

1 Introduction

Does short-term equity outperform long-term equity? Most asset pricing models imply risk premia and Sharpe ratios on short-term claims that are close to zero and thus much lower than those on long-term claims (Campbell and Cochrane 1999; Bansal and Yaron 2004). Yet, the empirical evidence on this topic is still debated. Van Binsbergen, Brandt, and Koijen (2012) (henceforth BBK) estimate prices of dividend strips with maturities close to two years from index options during 1996 through 2009 and show that these short-term dividend strips deliver higher risk premia and Sharpe ratios than does the long-term market. Van Binsbergen et al. (2013), Van Binsbergen and Koijen (2017), Cejnek and Randl (2020), and Gormsen (2021) extend this evidence using proprietary data on over-the-counter dividend swaps. Dividend swaps allow for direct measurement of dividend strip prices and are thus less susceptible to measurement error than strip prices estimated from options data (Boguth, Carlson, Fisher, and Simutin 2023). However, over-the-counter dividend derivatives started trading only in the early 2000s; exchange traded dividend futures started trading in the U.S. even later in 2015. Therefore, the use of dividend futures does not overcome the critique of Bansal, Miller, Song, and Yaron (2021) that short samples might be unrepresentative and tainted by an oversampling of recessions.

We come full circle by going back to the options data. We purchase the original data covering the BBK sample and extend the data until the end of 2022. We thus almost double the original BBK sample. The proportion of recessions in our sample is comparable to the historical occurrence of recessions.[1] Using public options data instead of proprietary dividend futures data allows us to share our data.[2]

[1] Bansal, Miller, Song, and Yaron (2021) show that the long run recession frequency in the U.S. is 15% and the frequency of severe recessions is 4%. We have National Bureau of Economic Research (NBER)-defined recessions in 9% of our sample.

[2] Giglio, Kelly, and Kozak (2020) and Gonçalves (2021a) estimate the equity term structure from the cross section of equity returns. That approach extends the data back even further than by using index options, but it comes at the cost

We take two steps to address potential measurement error in dividend strip prices. First, like BBK, we estimate dividend strip prices from time-matched put-call pairs. This approach is sensitive to the choice of the interest rate (Song 2016). Instead of using an exogenous interest rate that could bias results (BBK use the zero curve rate), we use an option-implied interest rate (Van Binsbergen, Diamond, and Grotteria 2022).

Second, we suggest the use of longer holding period logarithmic returns. Boguth, Carlson, Fisher, and Simutin (2023) advocate the use of logarithmic returns to remove the measurement error bias in average returns. We follow their advice throughout the paper. However, the use of logarithmic returns will not produce unbiased estimates of Sharpe ratios, because measurement error inflates return standard deviation (Blume and Stambaugh 1983) and, thus, biases Sharpe ratios downwards. To mitigate the effect of measurement error, we use logarithmic returns over longer holding periods. Theoretically, the measurement error vanishes asymptotically. Empirically, we find that it disappears for holding periods longer than two years.[3]

In presenting our results, we start with the unconditional results for the full sample. Like BBK, we fix the maturity of the strip at close to two years. We then measure term premia as the difference between market and strip returns for a given holding period (as in BBK and Gormsen 2021). We find the term premia to be rather flat or insignificantly upward-sloping. This compares to insignificantly downward-sloping term premia in BBK and insignificantly upward-sloping term premia in Bansal, Miller, Song, and Yaron (2021).

While term premia are mostly flat, we find the term structure of Sharpe ratios (the difference

of estimating the dynamics of the economy and investor preferences. Our data provide direct estimates of dividend strip prices and can be used as a yardstick to evaluate such alternative methods.

[3]Schulz (2016) challenges the BBK findings on the grounds of differential tax treatment of capital gains and dividends. It is unclear how differential tax treatment could explain why the volatility of strip returns decreases as we increase the holding period (see our results below). Moreover, Van Binsbergen and Koijen (2016) argue that the estimated tax rates of Schulz (2016) are unreasonably high. Using lower tax rates from the literature (Sialm 2009), they confirm the BBK findings. They also show that findings hold when comparing returns on dividend futures with different maturities, which are affected by taxes in the same way.

between market and strip Sharpe ratios) to be generally downward-sloping. For holding periods of two to three years, where the effect of the measurement error is mitigated, dividend strip volatility is substantially lower than market volatility, and the dividend strip Sharpe ratio is larger than the market Sharpe ratio, with typically significant slope. We always reject the hypothesis that the strip Sharpe ratio is zero. Our results thus compare to BBK and Van Binsbergen and Koijen (2017), except that we find that strip Sharpe ratios are high mainly due to their low volatility rather than due to their high returns. Results are robust to the choice of excess returns (in excess of the risk-free rate or in excess of duration-matched Treasury bond returns) and to using constant returns across holding periods.[4] We observe the same patterns in subsamples and when we consider methodological changes in estimating dividend strip returns (option maturity, interest rate, option moneyness, and transaction costs).

As an alternative to extending the holding period, we estimate a market model for dividend strip returns. Unlike Sharpe ratios, market model estimates are unaffected by the measurement error in strip returns. Indeed, we find that market model beta is low (as in BBK) and stable across holding periods (0.27, on average). Dividend strip alpha is positive (4% annualized and marginally significant) and also stable across holding periods. We interpret these results as further validation of our results and conclude that dividend strips perform well in terms of both Sharpe ratios and market alphas.

Next, we analyze how term premia vary over time. Guided by a simple present value model, we document that term premia are highly predictable by the market dividend-to-price ratio (henceforth market dp-ratio) and by the scaled difference between market and strip dp-ratios (henceforth scaled dp-ratios).[5] These results hold both in- and out-of-sample.

[4] Except that the term structure of Sharpe ratios turn insignificant — but stays downward-sloping — when using returns in excess of the risk-free rate in combination with constant means.

[5] Casella et al. (2023) link the time series variation in term premia to changes in investor optimism at long- and short-horizons. For research that uses information on dividend strips to predict market or strip returns as opposed to

Like Gormsen (2021), we interpret our predictive regressions as a statement about the business cycle (Campbell 1999). Our estimates of future realized term premia are positively associated with the current market dp-ratio. As bad times are characterized by a high market dp-ratio, our estimates suggest that term premia move countercyclically over the business cycle, in line with Gormsen (2021).[6] Using scaled dp-ratios or other measures of the business cycle leads to similar results.

We extend the predictive regression approach to predict the term structure of Sharpe ratios and market model alphas. We find that the difference in market and strip Sharpe ratios also moves countercyclically over the business cycle. In contrast, the difference in market and strip alphas moves procyclically over the business cycle.

All in all, after almost doubling the BBK sample period and after accounting for the effects of measurement error, we are able to confirm BBK's original result that dividend strips deliver high returns for their level of risk, both in terms of Sharpe ratios and market model alphas. We further confirm the result from Gormsen (2021) that term premia move countercyclically over the business cycle and provide new evidence on the business cycle variation in the term structure of Sharpe ratios and alphas.

We conclude by comparing our empirical findings to the predictions of theoretical asset pricing models. Among the models we consider, the 2-factor model of Gormsen (2021) seems to capture the highest number of our empirical findings. It produces a non-zero Sharpe ratio for short-term dividend strip returns and predicts that term premia move countercyclically over the business cycle. The habit model of Campbell and Cochrane (1999) and the long-run risk model of Bansal and

term premia, see BBK, Golez (2014), and Li and Wang (2018), among others.

[6]Van Binsbergen et al. (2013) and Bansal et al. (2021) focus on hold-to-maturity returns and document procyclical behavior. Gormsen (2021) shows that the two results are not inconsistent with one another since the term structure of hold-to-maturity returns varies both the maturity of the claims and the holding period, whereas we vary maturities of the claims (i.e., strip vs. market) for a given holding period (see also our discussion in Section 7).

Yaron (2004) align with our evidence regarding the countercyclicality of the term premia, but they predict close-to-zero strip Sharpe ratios, which we reject empirically. The value model of Lettau and Wachter (2007) and the rare disaster model of Gabaix (2012) generate positive strip Sharpe ratios, but they predict that term premia are either procyclical or constant across the business cycle, which we reject empirically.

2 Data

We obtain data on European S&P 500 index options (henceforth SPX options) from the Chicago Board of Options Exchange (CBOE). We use tick-level data for the period from January 1, 1990, through March 31, 2004, and minute-level data from January 1, 2004, through December 31, 2022. We aggregate the tick-level data to the minute level. The CBOE switched in the more recent data from Central Standard Time (Chicago) to Eastern Standard Time (New York City). We moved all time stamps to Central Standard Time. We merge the option data with the intradaily S&P 500 cash index from the Chicago Mercantile Exchange (CME). Data for long-maturity options in the early years are very sparse. Therefore, we follow BBK and start our analysis in January 1996. Our final time series is from January 1996 through December 2022.

We calculate realized dividends from the daily Datastream S&P 500 return index and the total return index. We use information on indicative dividends for the S&P 500 index from S&P Dow Jones Indices. We collect daily zero curve rates from January 1996 through December 2022 from OptionMetrics. We download nominal constant maturity Treasury interest rates from the H.15 filing of the St. Louis Federal Reserve Bank. We obtain returns on 2- and 10-year fixed maturity Treasuries from CRSP. We download the one-month Treasury bill rate and the market factor from Kenneth French's data library. For the business cycle analyses, we obtain the monthly series of annual consumption and the quarterly series of the output gap from the St. Louis Federal Reserve Bank and the monthly CAPE measure (market price over ten-year earnings) from Robert Shiller's website. For comparison, we also download the original data from BBK from the web page of the *American Economic Review*.

3 Methodology and Estimation

To study the attractiveness of dividend strips compared to the market, we need returns on both securities. We compute market returns from S&P 500 prices and dividends. We compute dividend strip prices from the put-call parity relation of European put and call options (SPX) on the S&P 500 index.

Put-call parity dictates that, at any given time t, the price of dividends on the underlying index during the life of the options P is given by:

$$P_t^\tau = S_t + p_t^\tau(X) - c_t^\tau(X) - Xe^{-rf_t^\tau \tau}, \tag{1}$$

where τ is the maturity of the options at time t, S is the value of the underlying index, p is the price of a European put option with strike price X, c is the price of a European call option with same strike price, and rf is the annualized continuously compounded risk-free rate of return over the corresponding period τ.

Van Binsbergen, Brandt, and Koijen (2012) estimate the price P of the short-term dividend strip using zero curve interest rates. Specifically, for a given day t and maturity τ, they find all intradaily pairs of put and call options with the same strike price and match them with the intradaily values of the index and the end-of-day values of the zero curve rate of the matching maturity. From each combination of the data with the same maturity, they estimate a strip price, which they aggregate into a single daily median price.

Results may be sensitive to the use of the zero curve interest rate. First, there is a time mismatch between end-of-day zero curve rates and intradaily data for the options and the index (Boguth, Carlson, Fisher, and Simutin 2023). Second, funding costs of marginal investors in index options may differ from the zero curve interest rate (Song 2016; Van Binsbergen and Koijen 2016). Ulrich, Florig, and Wuchte (2019) find that the unconditional term premia are either downward- or upward-sloping depending on which interest rate (OIS or LIBOR) they use as a proxy for the risk-free rate.

Even a small error in interest rates can lead to a large error in the estimated dividend strip returns. Interest rates that are too low (high) lead to strip prices that are also too low (high) (see Equation 1). Any mistake in the strip prices is then magnified in the calculation of strip returns. In the Internet Appendix A, we consider a simple calibration based on a small error of negative 6 basis points (bp). The error is based on our finding that option-implied interest rates at relevant maturities (as will be described below) are 6 bp higher than zero curve rates. We show that even such a small error can bias half-annual dividend strip returns by 0.86% (or 24.13% in relative terms). The elasticity of the strip return with respect to the interest rate error is large at −11.22.

3.1 Option-Implied Interest Rates

Therefore, we advocate the use of an interest rate invariant approach that relies on interest rates internally consistent with option prices. Specifically, we can treat Equation (1) as having two unknown variables, the dividend price P and the risk-free rate rf. Van Binsbergen, Diamond, and Grotteria (2022) identify the risk-free rate by combining two put-call parity relations with different strike prices X into a pair.[7]

[7] As an alternative, Golez (2014) identifies the risk-free rate by combining option data with futures data (see Demaskey and Heck (1998) for an early reference). Since standard SPX options expire on a monthly cycle and futures expire on a quarterly cycle, his approach restricts the set of possible maturities.

They discuss two different ways of estimating option-implied interest rates from such pairs. We refer to the first approach as the *outer product approach*. For a given maturity τ, we create all possible unique combinations of put-call pairs across strike prices. We denote the number of different put-call strike prices by N. The number of possible combinations is $A = \frac{N(N-1)}{2}$. For each put-call pair (indexed $a = 1,...A$), we compute an option-implied interest rate. That is, for each $i = 1,....N$ and for each $j = 1...N$, for which X_i is greater than X_j, we compute

$$rf_{t,\tau,a} = -\frac{1}{\tau} ln\left[\frac{\left(p_t^\tau(X_i) - c_t^\tau(X_i)\right) - \left(p_t^\tau\left(X_j\right) - c_t^\tau\left(X_j\right)\right)}{X_i - X_j}\right]. \tag{2}$$

Finally, we take the median implied rate as the daily option-implied interest rate. This approach is computationally intensive, but robust to outliers.

We refer to the second approach as the *regression approach*. For a given maturity τ, we run the following regression based on time-matched put-call parity relations:

$$S_t - c_t^\tau(X) + p_t^\tau(X) = P + \beta X + \epsilon. \tag{3}$$

We use the estimated coefficient for the strike price $\hat{\beta}$ to compute the implied risk-free rate, $rf = -\frac{1}{\tau} ln\left(\hat{\beta}\right)$. This is a computationally efficient method, but more sensitive to outliers than the outer product approach. Both methods produce the same estimates asymptotically.

We use the option-implied interest rates as an input in the put-call parity relation (Equation 1) to calculate prices of the dividend strips. Over the years, option trading has substantially increased. The data from the first part of the sample (1996 through 2003) are much sparser than from the second part (2004 through 2022). On January 31, 1996, and after filtering (see below), the number of unique option relations across all maturities is 3,271 (1,243 option relations for maturities greater than one year). On December 31, 2022, the number of unique option relations is 480,253 (87,941 option relations for maturities greater than one year). As a result, the impact of potential outliers is much larger in the first part of the sample, whereas computational speed is more of a concern in the latter part of the sample. We therefore use the outer product approach to estimate option-implied interest rates from 1996 through 2003 and the regression approach from 2004 through 2022. When we apply both approaches to a sample of recent data, we find that strip prices are virtually the same.[8]

Like BBK, we use options only on the last business day of each month between 10 a.m. and 2 p.m. We use standard monthly options that expire on the third Friday of each month. For the option price, we use the bid-ask midpoint and eliminate all options with bid or ask prices lower than $3. We also eliminate options with moneyness levels below 0.5 or above 1.5 and options with maturities of fewer than 90 days.

[Figure 1 about here]

We find that the one-year implied rate is on average 7 bp (2.82% in relative terms) higher than the zero curve rate. To illustrate, Figure 1 plots the one-year constant maturity implied rate along with the zero curve rate and the Treasury rate. We calculate constant maturity rates by linearly interpolating between the rates just below and above one year. The implied rate and the zero curve rate are substantially higher than the Treasury rate (by some 38 bp). The difference between the zero curve rate and the implied rate in the first half of the sample is rather small and amounts to 0.64% in relative terms. The difference between both rates increases in the second half of the sample to 11.80% in relative terms.[9] This suggests that the zero curve rate was a relatively good

[8] We also directly estimate the dividend prices in the outer product approach (substituting out the risk-free rate without estimating it) and in the regression approach (using $\hat{P}$ directly). In the early years of our sample period, the direct approaches lead to somewhat noisier dividend prices than the indirect approaches, which we prefer for that reason.

[9] One of the reasons the zero curve rate is lower than the option-implied interest rate in the recent sample is banks' underreporting of borrowing costs used in the calculation of LIBOR (Gandhi et al. 2019).

proxy for the funding costs of option investors in the earlier part of the sample (including the BBK sample), but that the use of the zero curve rate may overestimate dividend strip returns for recent years. We proceed to estimate dividend strip prices using our time- and maturity-matched option-implied interest rates. For comparison, we report estimations with the zero curve rate in the Internet Appendix B.2.

We follow BBK and estimate dividend strip prices using Equation (1). The only difference is that we use option-implied interest rates rather than zero curve rates. In estimating dividend strip prices, we use the same option pairs that we use in the calculation of option-implied interest rates. Each option pair gives us one estimate for the dividend price. We then take the median across all the dividend strip prices for a given maturity on a given day. For each month-end, we obtain estimates for dividend strip prices with maturities matching the option expiration dates.

Finally, we calculate monthly returns on dividend strips. We rely on the approach used by BBK and calculate returns from strip prices with maturities close to two years. To account for the availability and liquidity of longer-dated options, BBK focus on strip prices estimated from options expiring in either June or December.[10] Specifically, at the end of January of year t, we buy a dividend strip with a maturity of around 1.9 years (based on options expiring in December of year $t+1$). We then roll this strip for six months until we rebalance again at the end of July into a new dividend strip with 1.9 years to maturity (based on options expiring in June of year $t+2$). The only exception to this rule is July 2013 to January 2014, during which we let the strategy rely on the strip with maturity of around 1.5 years, because the appropriate options maturing in June 2015 were not listed until September 2013.

We then calculate the monthly strip return as the sum of the strip price at month-end plus the

[10]The liquidity of long-dated options is driven by the CBOE issuing cycle. For options expiring in more than three years, the CBOE initially lists December expirations, followed by June expirations. As we get closer to these expirations, the CBOE eventually adds other expiration months. For longer-dated options, liquidity is thus concentrated in the December and June expirations.

dividends that accrue in that month, divided by the strip price at the end of the previous month.[11] We refer to the resulting time-series as returns on short-term dividend strips.

We calculate monthly returns on the S&P 500 market so that we can compare them with strip returns. Throughout, we use logarithmic returns because they are less sensitive to the standard deviation bias and better estimate buy-and-hold returns accumulated over longer periods. We also define the market dp-ratio $\left(dp^{Mkt}\right)$ as the logarithm of the sum of dividends over the past year minus the logarithm of the current level of the S&P 500 index. For dividend strips, we define the strip dp-ratio $\left(dp^{Strip}\right)$ as the logarithm of the sum of dividends over the past year minus the logarithm of the price of a one-year dividend strip.

3.2 Measurement Error and Sharpe Ratios

We estimate strip returns from options. Any noise in the options data can lead to biased estimates of performance measures. Boguth, Carlson, Fisher, and Simutin (2023) argue that noise in the data may bias average returns upward and suggest using logarithmic returns. We follow their recommendation throughout the paper.

However, the use of logarithmic returns will not ensure unbiased Sharpe ratios. The reason is that noise will lead to negatively autocorrelated returns (Blume and Stambaugh 1983). Such negative autocorrelation inflates standard deviation estimates and lowers the Sharpe ratio. This effect is present even in the case of logarithmic returns.

To illustrate, we consider a simple model of monthly log prices with additive measurement error. The measurement error has variance σ_δ^2 that is uncorrelated with the log return variance σ_ε^2. See Internet Appendices C and particularly C.1 for details. The negative serial correlation induced by the measurement error inflates the variance of measured returns $var(\hat{r_{t,t+h}})$ relative to

[11] This is equivalent to reinvesting monthly dividends in dividend strips.

the variance of actual returns $var(r_{t,t+h})$ and hence depresses the measured Sharpe ratio relative to the actual Sharpe ratio. The effect dissipates with the length of the holding period. We can express the actual h-period Sharpe ratio $SR(r_t)$ as a function of the h-period measured Sharpe ratio $SR(\hat{r}_t)$:

$$SR(r_t) = \sqrt{\frac{var(\hat{r_{t,t+h}})}{var(r_{t,t+h})}} SR(\hat{r}_t) = \sqrt{1 + \left(\frac{2}{h}\right) \frac{\sigma_\delta^2}{\sigma_\varepsilon^2}} SR(\hat{r}_t). \tag{4}$$

For sufficiently long holding periods, the measured Sharpe ratio approaches the actual Sharpe ratio from below. In the empirical analysis, we consider holding periods of up to 36 months.

4 Unconditional Results

Table 1 reports the annualized summary statistics for single-period (monthly) logarithmic returns. Columns 1 and 2 presents statistics for raw market returns (8.54%) and raw dividend strip returns (7.10%). The difference in raw returns is insignificant.[12]

We provide two versions of excess returns. The first version uses market and strip returns in excess of the risk-free rate rf, where we use the one-month Treasury bill rate. Columns 3 and 4 report the statistics for market returns in excess of the risk-free rate (6.57%) and strip returns in excess of the risk-free rate (5.12%). The second version of excess returns matches the duration of the risk-free returns to the duration of the asset. The latter approach corresponds to returns on forward contracts and thus provides a comparison to studies that estimate dividend strip returns from dividend futures (Van Binsbergen et al. 2013; Van Binsbergen and Koijen 2017).[13] We

[12] Van Binsbergen, Brandt, and Koijen (2012) find that dividend strips offer insignificantly higher returns than the market. In the Internet Appendix D, we replicate their results and compare them to our estimates. We find that the difference stems from both our use of the option-implied interest rate and our extended sample. With the option-implied interest rate, dividend strips and the market deliver similar returns during the BBK sample, but dividend strips underperform during the extended sample (1996 through 2022).

[13] Since futures require no investment upfront, returns on dividend futures are already in excess of the returns on a bond with matching maturity.

subtract 2-year Treasury returns from dividend strip returns and 10-year Treasury returns from market returns.[14] Columns 5 and 6 give the market returns in excess of the 10-year Treasury rate (4.60%) and strip returns in excess of the 2-year Treasury rate (4.22%). The returns are closer to each other as the 10-year Treasury rate is higher than the 2-year Treasury rate.

[Table 1 about here]

[Figure 2 about here]

Figure 2 plots the cumulative returns for rolling over investments in the dividend strip or the market. $1.00 invested in the market for 27 years grows to $9.97, whereas $1.00 invested in the dividend strip grows to $6.76 (Panel A). Returns in excess of the risk-free rate (Panel B) grow to $5.86 and $3.97. Returns in excess of the Treasury rate (Panel C) grow to $3.45 and $3.11. Thus, accounting for the risk-free rate and, in particular, for the Treasury term structure matters for the comparison of market and strip returns, but neither the difference between the average returns nor the difference between the average excess returns are statistically significant.

Monthly dividend strip returns are approximately twice as volatile (32%) as monthly market returns (16%, 18% if we subtract Treasury returns). Monthly dividend strip returns also exhibit a strong negative autocorrelation of -0.33. This means that lagged strip returns explain 11% of the variation in monthly returns. In comparison, the AR(1) coefficient for the market return is close to zero at 0.02 (0.08 if we subtract Treasury returns).[15]

[14] Results are qualitatively similar if we subtract 20-year Treasury returns instead.

[15] When we check for a higher-order autocorrelation, we find that, of all AR coefficients out to six lags for both the strip and the market, only the AR(1) coefficient of the strip is significant.

4.1 Longer Holding Periods and Sharpe Ratios

A high standard deviation of single-period strip returns combined with strong negative serial correlation in returns is indicative of measurement error in dividend strip prices. In the Internet Appendix C.1, we show that the effect of measurement error should decline as the holding period increases over which we measure Sharpe ratios.

We consider holding periods of 1 through 36 months. That is, we sum the logarithmic returns for dividend strips and, separately, for the market over a given holding period: $r_t^h = \sum_{j=1}^{h} r_{t+1-j}$, for $h = 1, \ldots, 36$.[16] Figure 3 presents annualized standard deviations across different holding periods. Table 2 reports the corresponding summary statistics. Panels A always refer to returns in excess of the risk-free rate; Panels B refer to returns in excess of the Treasury rates.

[Table 2 about here]

[Figure 3 about here]

We note a drastic decrease in the annualized standard deviation for the dividend strip return in excess of the risk-free rate (Panel A) from 32% for monthly returns to 14% for annual returns before stabilizing at around 12% for holding periods beyond two years. This suggests that obtaining stable estimates for the standard deviation of dividend strip returns requires holding periods of at least two years. In comparison, the standard deviation for the market return in excess of the risk-free rate increases slightly from 16% (monthly) to 17% (annual) and then further to 18% (beyond two years). Overall, these patterns are consistent with a strong negative serial correlation for the dividend strip and a slightly positive serial correlation for the market. The patterns are very similar for returns in excess of Treasury bond returns (Panel B).

[Figure 4 about here]

[16] Note that, as we change the holding period, we do not change the maturity of the underlying asset.

These shifts in the standard deviation profoundly affect annualized Sharpe ratios (see Figure 4, Panel A for returns in excess of the risk-free rate). Specifically, the dividend strip Sharpe ratio increases from 0.16 for the monthly holding period to 0.48 for the three-year holding period. In contrast, the market Sharpe ratio decreases from 0.42 to 0.35. We test for the difference in Sharpe ratios using the heteroskedasticity- and autocorrelation-consistent (HAC) test proposed by Ledoit and Wolf (2008) using the asymptotically efficient QS kernel of Andrews (1991). We find that the strip Sharpe ratio is significantly different from the market Sharpe ratio for any holding period longer than 24 months. For holding periods beyond 6 months, we can also reject the null that the dividend strip Sharpe ratios are zero.[17] Since the measurement error is minimized at longer holding periods, we find that the dividend strip outperforms the market in terms of Sharpe ratios. Patterns for returns in excess of the Treasury bond returns are qualitatively similar, except that the Sharpe ratios are overall lower, see Panel B.

4.2 Longer Holding Periods and Mean Returns

One observation requires additional consideration. The documented changes in Sharpe ratios go beyond the changes in return volatility and are also due to changes in mean returns. For the market, the mean return in excess of the risk-free rate decreases from 6.57% to 6.22% as we increase the holding period. For the strip, the mean return in excess of the risk-free rate increases from 5.12% to 5.78%. These changes are driven by the fact that, for longer holding periods and when using overlapping observations, we place more weight on the observations from the center years of the sample period and less weight on early and late years. The mean return of each asset at different holding periods thus depends on how the asset performance at the center of the sample compares

[17] We test against zero by setting the market mean return equal to the risk-free rate rf, thus creating a hypothetical asset with zero Sharpe ratio and market volatility.

to the performance at the end points of the sample.

To isolate the effect of the changing mean, we recalculate Sharpe ratios by always using the mean of the monthly observations ($h = 1$).[18] We thus ensure that any changes in Sharpe ratios are driven by changes in the return volatility and not by changes in mean returns. We report the results for Sharpe ratios with constant mean along with the base case results in Table 2. The general pattern of Sharpe ratios with constant mean follows the pattern of the Sharpe ratios in the base case. There is a large increase in Sharpe ratios as we increase the holding period. In Panel A, the strip Sharpe ratio at the three year holding period is higher than the market Sharpe ratio, but the difference is small (0.37 vs 0.42) and not statistically significant.[19] Still, the strip Sharpe ratio with constant mean remains statistically significantly different from zero for longer holding periods. In Panel B, even the difference between the market Sharpe ratio and the strip Sharpe ratio with constant mean remains significant at longer holding periods.

As an alternative to isolating the effect of the changing mean, we wrap the data by connecting the last return with the first (the same approach is used in the circular bootstrap of Politis and Romano 1992). We then repeat our analysis using each month as a starting point. While unrealistic from the point of view of an investor, this approach ensures that the average return across the samples is the same regardless of the holding period. We repeat the main results from Table 2, Figure 3, and Figure 4 from the paper and report them in the Internet Appendix E. The general patterns are the same. The strip Sharpe ratio is sizable and statistically significantly different from zero. The strip Sharpe ratio is also higher than the market Sharpe ratio but insignificantly so.

[18] We thank an anonymous referee for this suggestion.

[19] We calculate p-values by demeaning at the given holding period and then adding back the constant mean ($h = 1$).

4.3 Longer Holding Periods and Robustness

The Internet Appendix B reports robustness results concerning (i) the choice of the dividend strip maturity, (ii) the use of the zero curve interest rate instead of the option-implied interest rate, (iii) the choice of option moneyness, and (iv) the effect of transactions costs by holding the strip to maturity instead of monthly rebalancing. Results are very similar to the results in the base case in that the strip Sharpe ratio always increases with the holding period. The strip Sharpe ratios are also comparable to the base case. The exception is case (ii). When using the zero curve interest rate instead of the option-implied interest rate, we overestimate strip returns and, thus, strip Sharpe ratios are overall higher.

4.4 Longer Holding Periods and Sharpe Ratios: Subsamples

Figure 5 plots dividend strip Sharpe ratios and market Sharpe ratios for different subsamples. Panel A plots the annualized Sharpe ratios using returns in excess of the risk-free rate. The left panel shows that strip Sharpe ratios are significantly higher than market Sharpe ratios at longer holding periods during the BBK period from January 1996 through October 2009, see the corresponding Panel A in Table F1 in the Internet Appendix F. The right panel shows that strip Sharpe ratios are similar to market Sharpe ratios at longer holding periods for the period from December 2004 through December 2022 (in Bansal et al. (2021), the time period is from December 2004 to February 2017), see the corresponding Panel A in Table F2. We note that the null of many leading asset pricing models (e.g., Campbell and Cochrane 1999 and Bansal and Yaron 2004) is that the Sharpe ratio of the short-term dividend strip should be close to zero, which we strongly reject. Panel B of Figure 5 plots the annualized Sharpe ratios for the market and the strip for returns in excess of Treasury bond returns, see Panels B in Tables F1 and F2 in the Internet Appendix F. Here, the strip

Sharpe ratio is always significantly higher than the market Sharpe ratio at longer holding periods.

[Figure 5 about here]

We are concerned about ending the first subsample in the middle of the global financial crisis. We thus let the first subsample run from January 1996 to November 2007. We report results in the Internet Appendix F in Table F3 and Figure F1 (left). Excluding the global financial crisis increases mean returns, yet the general patterns remain the same as in the base case. The strip Sharpe ratio is always higher than the market Sharpe ratio at longer holding periods.

We are also concerned about ending the second subsample after the Covid-related recession. We thus let the second subsample run from December 2004 to January 2020. We report results in the Internet Appendix F in Table F4 and Figure F1 (right). The general patterns remain the same as in the base case. The strip Sharpe ratio is always higher than the market Sharpe ratio at longer holding periods.

4.5 Longer Holding Periods and the Market Model

So far, we have focused on Sharpe ratios, which are important for the pricing of the market. Now, we estimate betas of dividend strips from a market model, which are important for pricing in the cross-section of stocks (Van Binsbergen, Brandt, and Koijen 2012; Gonçalves 2021b; Gormsen and Lazarus 2023). An added benefit of the market model is that the estimates of market risk are unbiased in the presence of measurement error. Specifically, we can write the market model for strip returns as:

$$r_{t,t+h}^{\widehat{Strip}} - rf_{t,t+h} = \alpha_{t,t+h} + \beta_{t,t+h}(r_{t,t+h}^{Mkt} - rf) + error_t, \tag{5}$$

where rf is the constant risk-free rate, $r_{t,t+h}^{\hat{Strip}}$ is the h-period strip return, and $r_{t,t+h}^{Mkt}$ the h-period market return. We use the hat symbol $\hat{Strip}$ to denote that monthly strip prices are observed with measurement error. Market prices do not contain measurement error. All returns are in logarithms to ensure that mean returns are not affected by measurement error. Since the measurement error does not affect the independent variable, the beta estimates are also unbiased (the Internet Appendix C.2 provides the details).[20] We thus expect betas to be relatively stable across the different holding periods. While measurement error does not introduce a bias in the estimated coefficients, it does affect the statistical significance of these estimates, even in the case of non-overlapping observations. We show in the Internet Appendix C.2 that confidence intervals around beta are wider in the presence of measurement error when compared to standard OLS confidence intervals of estimates without measurement error. We use Newey-West standard errors with 6 lags to correct for the autocorrelation in error terms induced by the measurement error. For longer holding periods, when h exceeds 6, we set the number of lags to h.

[Table 3 about here]

We use the S&P 500 index for the aggregate market and the one-month T-bill rate for the risk-free rate. Consistent with Van Binsbergen, Brandt, and Koijen (2012), we find that the estimated betas are relatively low (at around 0.27), see Table 3. We also find that betas are rather stable across holding periods. Beta starts at 0.32 at the monthly holding period, then decreases to 0.24 at the 24-month holding period, and increases again to 0.29 at the 36-month holding period. It is statistically significant at the 10% level at short (up to 18 months) and very long (36 months) holding periods, while it is insignificant at intermediate holding periods. The estimated annualized alpha is 4% across holding periods (3% at the monthly horizon) and significant at the 10% level

[20] We thank an anonymous referee for this suggestion and the argument that market estimates are unbiased in the presence of measurement error when mean returns are unbiased.

for holding periods of 12 months and longer.

We conduct two robustness checks. First, using a broader proxy for the market from the Fama-French data library, we verify in Table F5 in the Internet Appendix F that results are similar to the main results. Second, in line with the literature standard, we repeat our market model using raw returns (not in logarithms).[21] Results in Table F6 in the Internet Appendix F are very similar to the main results. The exception is the alpha for short holding periods, which is upward biased (8% at the monthly holding period).

The fact that market model betas for dividend strips are low and stable across horizons is consistent with our argument that the high strip volatility for short holding periods is due to measurement error. In fact, if strip volatility were driven by systematic discount rate fluctuations, a strip volatility higher than the market volatility would necessitate betas above one at short holding periods. Still, we cannot rule out the possibility that the strip volatility at short holding periods could be driven by its own unique discount-rate variation.

5 Predictability of Realized Term Premia

In this section, we investigate whether realized term premia (i.e., market returns minus strip returns) are predictable by dividend-to-price ratios. There is a long tradition of predicting market returns using the market dp-ratio (Campbell and Shiller 1988). Gormsen (2021) analyzes how term premia vary with the market dp-ratio and the ratio of strip and market prices. Among other specifications, Van Binsbergen, Brandt, and Koijen (2012) use the strip dp-ratio to predict returns on dividend strips. We relate term premia to the market dp-ratio and to the scaled difference between the market and strip dividend-to-price ratios (i.e., scaled dp-ratios). We motivate our analysis

[21] We used logarithmic returns in the base case to preserve the mean return in the presence of measurement error.

with a simple present value model.

5.1 Present Value Model

We assume that the market is infinitely lived, whereas the dividend strip lives only for one period. They both share the same dividend growth process. We assume expected changes in log dividends $E(\Delta d_{t+1}) = g_t$ to be an AR(1) process:

$$g_t = \gamma_0 + \gamma_1 g_{t-1} + \varepsilon_t^g, \tag{6}$$

We allow the expected returns on the market to follow a different process from the expected returns on the dividend strip. Specifically, we assume that expected market returns $E\left(r_{t+1}^{Mkt}\right) = \mu_t^{Mkt}$ follow an AR(1) process:

$$\mu_t^{Mkt} = \delta_0 + \delta_1 \mu_{t-1}^{Mkt} + \varepsilon_t^{\mu,Mkt}. \tag{7}$$

For expected strip returns, we do not impose any specific dynamics, $E\left(r_{t+1}^{Strip}\right) = \mu_t^{Strip}$. Under these assumptions and using the Campbell and Shiller (1988) decomposition, we can write the logarithm of the market dp-ratio as (Van Binsbergen and Koijen 2010):

$$dp_t^{Mkt} = \kappa + \left(\frac{1}{1-\rho\delta_1}\right)\mu_t^{Mkt} - \left(\frac{1}{1-\rho\gamma_1}\right)g_t, \tag{8}$$

where $\rho = \frac{exp(-\bar{dp})}{1+exp(-\bar{dp})}$. The logarithm of the strip dp-ratio is $dp_t^{Strip} = \mu_t^{Strip} - g_t$. The difference between expected market and strip returns is:

$$E\left(r_{t+1}^{Mkt}\right) - E\left(r_{t+1}^{Strip}\right) = -A\kappa + \left[A dp_t^{Mkt} - dp_t^{Strip}\right] + B g_t, \tag{9}$$

where $A = (1-\rho\delta_1)$, $B = \left(\frac{\rho\gamma_1 - \rho\delta_1}{1-\rho\gamma_1}\right)$. Thus, the term structure of expected returns is related to the variation in both the scaled dp-ratios and the expected growth rate.

5.2 In-Sample Predictability

We use Equation (9) to motivate our predictive regressions. We regress differences in market and strip returns (i.e., term premia) over the next h periods on predictor variables X_t:

$$\sum_{j=1}^{h}\left(r_{t+j}^{Mkt} - r_{t+j}^{Strip}\right) = \alpha + \beta X_t + \varepsilon_t, \tag{10}$$

where $h = 12, .., 36$ months. All variables are in logarithms. For the predictor variables X_t, we use the market dp-ratio (dp_t^{Mkt}) and the scaled dp-ratios ($A^h dp_t^{Mkt} - dp_t^{Strip}$, where the scaling factor $A^h = \left(1-\rho\delta_1^h\right)$ depends on ρ, δ_1, and h). During our sample period, the coefficient ρ is 0.9824. To estimate the persistence of expected returns, we follow Golez and Koudijs (2023) and infer the persistence of expected returns $\theta(h)$ from a predictive regression of market returns over the next h months on lagged values of the market dp-ratio as $\theta(h) = \frac{1-\delta_1^{h/12}}{1-\delta_1}\,\theta(12)$. We estimate $\hat{\theta}(36) = 0.91$ and $\hat{\theta}(12) = 0.35$, which implies an annual persistence of expected returns of $\delta_1 = 0.85$. We are concerned about measurement error in strip prices feeding into the strip dp-ratio and causing spurious predictability.[22] As in Van Binsbergen, Brandt, and Koijen (2012), we therefore average the strip dp-ratios of the last three months.[23]

We also consider adding a proxy for expected dividend growth to our set of predictors. Specifically, we use the logarithm of indicated dividends over the 12-month trailing sum of dividends. The indicated dividends over the next year are provided by S&P Dow Jones Indices. They are based

[22] We thank an anonymous referee for pointing out this concern.

[23] As expected, our predictability results are stronger when we do not average the strip dp-ratios, see Table F7 in the Internet Appendix F.

on announced dividends, or, if dividends have not yet been announced, they are the last announced dividends projected into the future.

We report the results in Table 4. Newey-West t-statistics appear in parentheses, with the number of lags equal to h. In brackets are the average t-statistics from h non-overlapping samples starting in months $1, 2, ..., h$. For each non-overlapping sample, we estimate the Newey-West t-statistic with three lags (to account for any residual effect of the measurement error on standard errors).

[Table 4 about here]

The market dp-ratio is positively associated with future realized term premia. It best predicts the realized term premia over longer holding periods (coefficients are statistically significant at the 5% level for holding periods longer than 18 months). The estimated coefficient on the scaled dp-ratios is also positive and significant at any holding period. The R-squared varies from 17% at the annual holding period to 46% at the three-year holding period.[24] When we add indicated dividend growth as an additional predictor, the R-squared increases to 24% at the annual holding period and to 49% at the three-year holding period. The estimated coefficient for the indicated dividend growth is negative, which is consistent with the notion that expected returns are more persistent than expected growth rates (Golez and Koudijs 2023).[25]

6 Conditional Results

This section analyzes how term premia, Sharpe ratios, and alphas vary over the business cycle.

[24] The scaled dp-ratios also predict realized term premia out-of-sample, see Internet Appendix G.

[25] In Equation (9), the coefficient B is negative iff $\delta_1 > \gamma_1$.

6.1 Term Premia over the Business Cycle

Gormsen (2021) argues that the predictive regression of term premia on the current value of the market dp-ratio reveals how term premia vary over the business cycle. The idea goes back to Campbell (1999), who argued that a high dp-ratio reflects high risk premia that arise due to a high price of risk. A high market dp-ratio thus signals bad times, whereas a low market dp-ratio signals good times. Gormsen (2021) finds a positive relation between the market dp-ratio and the term premia measured over the next 12 months and interprets that as evidence for term premia that move countercyclically over the business cycle. Above in Table 4, Panel A, the term premia measured over the next 12 to 36 months are positively related to the current value of the market dp-ratio. The relation is even stronger if we use the scaled dp-ratios (Panel B). Our evidence is thus consistent with Gormsen (2021) and suggests that term premia move countercyclically over the business cycle.[26] In Table F8 in the Internet Appendix F, we use alternative measures of good and bad times. We use the market ten-year earnings-to-price ratio (the inverse of Robert Shiller's Cyclically Adjusted Price Earnings (CAPE) ratio), the change in consumption, and the output gap. While the results are not as strong as those with the dp-ratios, they all point in the direction of countercyclical term premia.

6.2 Sharpe Ratios over the Business Cycle

Next, we test how Sharpe ratios vary over good and bad times. We consider the same predictive regressions as before, except that we replace the term premia with the difference between market

[26]Our measure of term premia (i.e., the difference between market and strip returns) is the same as in Gormsen (2021), and, thus, our results are comparable. In contrast, Van Binsbergen et al. (2013) and Bansal et al. (2021) focus on hold-to-maturity returns and document procyclical behavior. Gormsen (2021) shows that the two results are not inconsistent with one another since the term structure of hold-to-maturity returns varies both the maturity of the claims and the holding period, whereas we vary maturities of the claims (i.e., strip vs. market) for a given holding period (see also our discussion in Section 7).

and strip Sharpe ratios:

$$SR^{Mkt}_{t,t+h} - SR^{Strip}_{t,t+h} = \alpha + \beta X_t + \varepsilon_t, \tag{11}$$

where $h = 12,..,36$ months. We calculate Sharpe ratios as the difference between the mean returns in excess of the risk-free rate and the volatility of excess returns over the next h months. For the predictors, we use either the market dp-ratio (dp^{Mkt}_t) or the scaled dp-ratios ($A^h dp^{Mkt}_t - dp^{Strip}_t$).

Results are reported in Table 5. We find that the term structure of Sharpe ratios also moves countercyclically over the business cycle and, based on the market dp-ratio, significantly so at the 5% level out to 30 months. As with the term premia, the difference in Sharpe ratios is even more strongly related to the scaled dp-ratios and significantly so at all holding periods. In the Internet Appendix F, we confirm that results are similar if we use returns in excess of Treasury returns (Table F9) and qualitatively the same when we use alternative measures for good and bad times (Table F10).

[Table 5 about here]

6.3 Market Model Alphas over the Business Cycle

Finally, we test how market model alphas vary over good and bad times. We consider the same predictive regressions as before, except that we replace the term premia with the difference in market and strip alpha (where market alpha is zero by construction):

$$\alpha^{Mkt}_{t,t+h} - \alpha^{Strip}_{t,t+h} = -\alpha^{Strip}_{t,t+h} = \alpha + \beta X_t + \varepsilon_t, \tag{12}$$

where $h = 12,..,36$ months. We calculate strip alphas according to Eq. (5) over the next h months. For the predictors, we use either the market dp-ratio (dp^{Mkt}_t) or the scaled dp-ratios ($A^h dp^{Mkt}_t -$

dp_t^{Strip}).

Results are reported in Table 6. We find that the term structure of alphas is negatively related to the dp-ratios and thus moves procyclically over the business cycle. The relationship is strongest for longer holding periods. In the Internet Appendix F, we confirm that the results are qualitatively the same when we use alternative measures for good and bad times (Table F11).

[Table 6 about here]

7 Discussion of Results

We collect our main empirical findings in Table 7. We investigate three economic aspects of equity returns: (i) term premia, (ii) the term structure of Sharpe ratios, and (iii) the term structure of market model alphas. For each economic aspect, we analyze the unconditional slope and the movement over the business cycle. We also note in Table 7, which of these economic aspects have been considered in the literature already.

[Table 7 about here]

The literature uses different term structures. We define term premia as the difference between market and strip returns for a given holding period. That is, for a given holding period, we only vary the maturity of the asset. This is similar to BBK and Gormsen (2021) but different from Van Binsbergen et al. (2013), Van Binsbergen and Koijen (2017), and Bansal et al. (2021), who focus on hold-to-maturity returns and compare returns on assets with different maturities over their respective maturities. Hence, they vary both the maturity of the assets and the holding period.[27]

[27] Gormsen (2021) provides a detailed discussion of these differences and shows that they are important (e.g., countercyclical term premia that vary only the maturity of the asset can be consistent with procyclical term premia of hold-to-maturity returns).

To ensure comparability, we mainly relate our empirical results to those reported in BBK and Gormsen (2021).

Our main findings can be summarized as follows. We find that strips with close to two years maturity deliver significant Sharpe ratios and alphas. While term premia are rather flat and insignificantly sloped, Sharpe ratios and alphas are significantly downward-sloping. Over the business cycle, term premia move countercyclically, as do differences in Sharpe ratios, whereas differences in alphas move procyclically. The latter two aspects are new to the literature.

We next assess five theoretical models and their predictions concerning our empirical findings. We gather the theoretical predictions of these models from BBK (in particular Figures 5 and 6 of BBK), Table 1 from Gormsen (2021), and Van Binsbergen and Koijen (2017). For these models, we mostly have slope predictions for term premia (and also some level predictions for strip returns), Sharpe ratios, and alphas while we have few cyclicality predictions.

7.1 Bansal and Yaron (2004) long-run risk and Campbell and Cochrane (1999) habit formation models

Both the long-run risk and the habit formation models fail to capture unconditional quantities. According to the models, strip excess returns and, thus, the strip Sharpe ratio are close to zero, while, empirically, they are significantly positive. Moreover, the models predict an upward-sloping term structure of Sharpe ratios, while, empirically, it is downward-sloping. However, both models do predict the countercyclicality of term premia, which is borne out empirically.

7.2 Gabaix (2012) rare disaster model

The rare disaster model predicts flat term premia and a downward-sloping term structure of Sharpe ratios, consistent with the data. Yet it also predicts that term premia do not vary over the business cycle, while, empirically, they move countercyclically.

7.3 Lettau and Wachter (2007) value premium model

The value premium model predicts non-negligible premia for short-term strips, consistent with the data. However, it predicts downward-sloping instead of the rather flat term premia. Moreover, it predicts term premia to move procyclically over the business cycle as opposed to our empirical finding of countercyclical movement. The predicted downward-sloping term structure of Sharpe ratios is consistent with the data.

7.4 Gormsen (2021) 2-factor model

The final model we consider is the 2-factor model of Gormsen (2021). The model predicts downward-sloping instead of the rather flat term premia that we find in the data. But, consistent with our empirical evidence, it predicts non-negligible premia for short term strips, and it predicts term premia to move countercyilically over the business cycle.

We conclude that different models capture different aspects of the equity term structure but no model fits them all. We hope that our empirical findings can help guide the development of new theoretical models.

8 Concluding Remarks

We estimate dividend strip prices from intradaily data for options on the S&P 500 index from 1996 through 2022. We almost double the existing time series of strip prices, use an option-implied interest rate to avoid biases because of the use of exogenous interest rates, and advocate the use of longer holding period returns to minimize the effect of measurement error in dividend strip prices. We show that strip Sharpe ratios are at least as large as market Sharpe ratios, as long as we focus on longer holding periods, where the effect of the measurement error is marginal. These results hold when we use returns in excess of the short-term risk-free rate or when we match the duration of equity and bond returns. The Sharpe ratio results are matched by a significantly positive strip alpha from a market model. Over the business cycle, the term premia and the term structure of Sharpe ratios move countercyclically, whereas the term structure of alphas moves procyclically. Overall, our results are most consistent with the predictions of the 2-factor model of Gormsen (2021), which predicts a downward-sloping term structure for Sharpe ratios and term premia that move countercyclically over the business cycle.

The replication code and data are available in the Harvard Dataverse at:

Jackwerth, Jens, 2023, "Replication Data for: Holding Period Effects in Dividend Strip Returns", https://doi.org/10.7910/DVN/D5LGGT, Harvard Dataverse, DRAFT VERSION

9 References

Andrews, D. W. K. 1991. Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica* 59:817–858.

Bansal, R., S. Miller, D. Song, and A. Yaron 2021. The term structure of equity risk premia. *Journal of Financial Economics* 142:1209–1228.

Bansal, R. and A. Yaron 2004. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59:1481–1509.

Blume, M. E. and R. F. Stambaugh 1983. Biases in computed returns: An application to the size effect. *Journal of Financial Economics* 12:387–404.

Boguth, O., M. Carlson, A. Fisher, and M. Simutin 2023. The term structure of equity risk premia: Levered noise and new estimates. *Review of Finance* 27:1155–1182.

Campbell, J. Y. 1999. Asset prices, consumption, and the business cycle. In *Handbook of macroeconomics vol. 1*, ed. J. B. Taylor and M. Woodford, 1231–1303. Amsterdam, Netherlands: Elsevier.

Campbell, J. Y. and J. H. Cochrane 1999. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107:205–251.

Campbell, J. Y. and R. J. Shiller 1988. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies* 1:195–228.

Cassella, S., B. Golez, H. Gulen, and P. Kelly 2023. Horizon bias and the term structure of equity returns. *Review of Financial Studies* 36:1253–1288.

Cejnek, G. and O. Randl 2020. Dividend risk premia. *Journal of Financial and Quantitative Analysis* 55:1199–1242.

Demaskey, A. L. and J. L. Heck 1998. Put-call-forward exchange parity and the implied risk-free rate. In *Advances in investment analysis and portfolio management vol 5*, ed. C.-F. Lee, 70–85. Oxford, UK: JAI Press.

Gabaix, X. 2012. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics* 127:645–700.

Gandhi, P., B. Golez, J. C. Jackwerth, and A. Plazzi 2019. Financial market misconduct and public enforcement: The case of Libor manipulation. *Management Science* 65:5268–5289.

Giglio, S., B. T. Kelly, and S. Kozak 2020. Equity term structures without dividend strips data. Working paper, Yale University.

Golez, B. 2014. Expected returns and dividend growth rates implied by derivative markets. *Review of Financial Studies* 27:790–822.

Golez, B. and P. Koudijs 2023. Equity duration and predictability. Working paper, University of Notre Dame.

Gonçalves, A. S. 2021a. Reinvestment risk and the equity term structure. *Journal of Finance* 76:2153–2197.

Gonçalves, A. S. 2021b. The short duration premium. *Journal of Financial Economics* 141:919–945.

Gormsen, N. J. 2021. Time variation of the equity term structure. *Journal of Finance* 76:1959–1999.

Gormsen, N. J. and E. Lazarus 2023. Duration-driven returns. *Journal of Finance* 78:1393–1447.

Ledoit, O. and M. Wolf 2008. Robust performance hypothesis testing with the Sharpe ratio. *Journal of Empirical Finance* 15:850–859.

Lettau, M. and J. A. Wachter 2007. Why is long-horizon equity less risky? A duration-based explanation of the value premium. *Journal of Finance* 62:55–92.

Li, Y. and C. Wang 2018. Rediscover predictability: Information from the relative prices of long-term and short-term dividends. Working Paper, University of Washington.

Politis, D. N. and J. P. Romano 1992. A circular block-resampling procedure for stationary data. In *Exploring the limits of bootstrap*, eds. R. LePage and L. Billard, 263–270. New York, NY: John Wiley.

Schulz, F. 2016. On the timing and pricing of dividends: Comment. *American Economic Review* 106:3185–3223.

Sialm, C. 2009. Tax changes and asset pricing. *American Economic Review* 99:1356– 83.

Song, Y. 2016. Dealer funding costs: Implications for the term structure of dividend risk premia. Working paper, University of Washington.

Ulrich, M., S. Florig, and C. Wuchte 2019. A model-free term structure of US dividend premiums. Working paper, Karlsruhe Institute of Technology.

Van Binsbergen, J., M. Brandt, and R. Koijen 2012. On the timing and pricing of dividends. *American Economic Review* 102:1596–1618.

Van Binsbergen, J., W. Hueskes, R. Koijen, and E. Vrugt 2013. Equity yields. *Journal of Financial Economics* 110:503–519.

Van Binsbergen, J. and R. Koijen 2010. Predictive regressions: A present-value approach. *Journal of Finance* 65:1439–1471.

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

Van Binsbergen, J., W. F. Diamond, and M. Grotteria 2022. Risk-free interest rates. *Journal of Financial Economics* 143:1–29.

Van Binsbergen, J. and R. Koijen 2016. On the timing and pricing of dividends: Reply. *American Economic Review* 106:3224–3237.

Tables and Figures

Table 1: **Monthly Returns (Annualized)**

	Market	Strip	Market	Strip	Market	Strip
	Log Return		Minus rf		Minus Treasury ret.	
Mean	8.54%	7.10%	6.57%	5.12%	4.60%	4.22%
Std. dev.	15.68%	31.98%	15.71%	31.99%	18.08%	31.98%
Sharpe ratio			0.42	0.16	0.25	0.13
AR(1)	0.02	-0.33	0.02	-0.33	0.08	-0.33
N	323	323	323	323	323	323

Table 1 presents summary statistics for the monthly returns. Returns are continuously compounded (in logarithms), annualized, and expressed as a percentage. Risk-free rate (rf) is the one-month Treasury bill rate. For returns in excess of Treasury returns, we subtract 10-year Treasury bond returns from the market returns and 2-year Treasury bond returns from the strip returns. The period is from January 1996 through December 2022.

Table 2: **Holding Period Returns and Sharpe Ratios (Annualized)**

	1m	6m	12m	18m	24m	30m	36m
Panel A: Returns in excess of the risk-free rate							
Market-ret. - rf							
Mean	6.57%	6.67%	6.85%	6.91%	6.75%	6.48%	6.22%
Std. dev.	15.71%	16.41%	17.29%	17.75%	18.08%	17.85%	17.68%
Sharpe ratio	0.42	0.41	0.40	0.39	0.37	0.36	0.35
Sharpe ratio (Const. mean)	0.42	0.40	0.38	0.37	0.36	0.37	0.37
Strip ret. - rf							
Mean	5.12%	5.71%	5.72%	5.89%	5.91%	5.88%	5.78%
Std. dev.	31.99%	18.48%	14.39%	13.22%	12.46%	12.09%	12.10%
Sharpe ratio	0.16	0.31	0.40	0.45	0.47	0.49	0.48
Diff. (*p*-val.)	(0.21)	(0.62)	(0.99)	(0.49)	(0.06)	(0.00)	(0.01)
Diff. wrt. zero [*p*-val.]	[0.40]	[0.06]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
Sharpe ratio (Const. mean)	0.16	0.28	0.36	0.39	0.41	0.42	0.42
Diff. (*p*-val.)	(0.21)	(0.53)	(0.84)	(0.83)	(0.39)	(0.23)	(0.32)
Diff. wrt. zero [*p*-val.]	[0.40]	[0.10]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
N	323	318	312	306	300	294	288
Panel B: Returns in excess of the Treasury bond returns							
Market ret. - 10y Treasury ret.							
Mean	4.60%	4.38%	4.35%	4.18%	3.94%	3.58%	3.29%
Std. dev.	18.08%	19.60%	20.00%	20.51%	20.77%	20.10%	19.56%
Sharpe ratio	0.25	0.22	0.22	0.20	0.19	0.18	0.17
Sharpe ratio (Const. mean)	0.25	0.23	0.23	0.22	0.22	0.23	0.24
Strip ret. - 2y Treasury ret.							
Mean	4.22%	4.71%	4.67%	4.77%	4.76%	4.70%	4.59%
Std. dev.	31.98%	18.78%	14.72%	13.65%	13.06%	12.78%	12.78%
Sharpe ratio	0.13	0.25	0.32	0.35	0.36	0.37	0.36
Diff. (*p*-val.)	(0.55)	(0.87)	(0.34)	(0.01)	(0.00)	(0.00)	(0.00)
Diff. wrt. zero [*p*-val.]	[0.50]	[0.09]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
Sharpe ratio (Const. mean)	0.13	0.22	0.29	0.31	0.32	0.33	0.33
Diff. (*p*-val.)	(0.55)	(0.95)	(0.59)	(0.16)	(0.00)	(0.01)	(0.01)
Diff. wrt. zero [*p*-val.]	[0.50]	[0.13]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
N	323	318	312	306	300	294	288

Table 2 presents summary statistics for the holding period returns ranging from $h = 1$ month through $h = 36$ months. Returns are continuously compounded (in logarithms), annualized, and expressed as a percentage. Risk-free rate (rf) is the one-month Treasury bill rate. Panel A reports results for returns in excess of the risk-free rate. Panel B reports results for returns in excess of the Treasury bond returns. Sharpe ratio for a holding period h is the mean excess return over the standard deviation of excess returns, measured for the same holding period h. Sharpe ratio (Const. mean) is calculated as the mean excess return at the monthly holding period $h = 1$ over the standard deviation of excess returns at a given holding period h. In parentheses are p-values for the HAC test of Ledoit and Wolf (2008) for the difference in Sharpe ratios between excess returns for the strip and the excess returns for the market. In brackets are the p-values for the same HAC test based on demeaned excess market returns. The period is from January 1996 through December 2022.

Table 3: **Holding Period Returns and the Market Model (Annualized Returns)**

	1m	6m	12m	18m	24m	30m	36m
$alpha$	0.03	0.04	0.04	0.04	0.04	0.04	0.04
$t-stat.$	(0.81)	(1.46)	(1.79)	(1.87)	(1.85)	(1.84)	(1.77)
$beta$	0.32	0.29	0.27	0.24	0.24	0.26	0.29
$t-stat.$	(1.80)	(2.19)	(2.10)	(1.65)	(1.52)	(1.55)	(1.68)
R^2	0.03	0.07	0.11	0.10	0.12	0.14	0.18
N	323	318	312	306	300	294	288

Table 3 presents market model estimates for dividend strip returns in excess of the risk-free rate for holding periods ranging from $h = 1$ month through $h = 36$ months. Returns are continuously compounded (in logarithms), annualized, and expressed as a percentage. We use the S&P 500 index as a proxy for the market. The risk-free rate is the one-month Treasury bill rate. In parentheses are the Newey-West t-statistics with $max(h,6)$ lags, where h is the holding period expressed in months. The period is from January 1996 through December 2022.

Table 4: **Predicting the Realized Term Premia**

	12m	18m	24m	30m	36m
Panel A:					
dp_t^{Mkt}	0.28	0.45	0.65	0.76	0.86
$t-stat(Overlap.)$	(1.83)	(2.04)	(2.53)	(2.66)	(2.85)
$t-stat(Nonoverlap.)$	[1.91]	[2.41]	[3.52]	[3.36]	[3.09]
R^2	0.10	0.16	0.25	0.31	0.37
Panel B:					
$A^h dp_t^{Mkt} - dp_t^{Strip}$	0.69	0.95	1.21	1.34	1.46
$t-stat(Overlap.)$	(2.89)	(2.83)	(3.42)	(4.25)	(5.84)
$t-stat(Nonoverlap.)$	[2.47]	[2.31]	[3.53]	[3.90]	[4.73]
R^2	0.17	0.24	0.33	0.39	0.46
Panel C:					
$A^h dp_t^{Mkt} - dp_t^{Strip}$	1.04	1.36	1.56	1.53	1.64
$t-stat(Overlap.)$	(4.74)	(4.46)	(5.08)	(5.38)	(7.20)
$t-stat(Nonoverlap.)$	[3.42]	[3.42]	[4.37]	[3.67]	[4.36]
g_t^{Ind}	-1.35	-1.74	-1.68	-1.08	-1.13
$t-stat(Overlap.)$	(-5.13)	(-7.10)	(-7.00)	(-4.31)	(-4.04)
$t-stat(Nonoverlap.)$	[-2.85]	[-2.80]	[-3.33]	[-1.82]	[-1.56]
R^2	0.24	0.33	0.40	0.41	0.49
N	310	304	298	292	286

Table 4 presents the results of the predictive regressions for the difference between market and strip returns for holding periods ranging from $h = 12$ month through $h = 36$ months. In panel A, the predictor variable is the market dp-ratio. In Panel B, the predictor variable is the scaled difference between the market and strip dp-ratios. In panel C, we add a control for indicative dividend growth. We take the average for the strip dp-ratios over the last three months. In parentheses are t-statistics based on the Newey-West (1987) correction with h lags. In brackets are the average t-statistics from h non-overlapping samples starting in months $1, \ldots, h$. For each non-overlapping sample, we set the number of lags for the Newey-West correction equal to 3. N reports the number of overlapping observations. The period is from January 1996 through December 2022.

Table 5: **Predicting the Term Structure in Sharpe Ratios**

	12m	18m	24m	30m	36m
Panel A:					
dp_t^{Mkt}	7.30	10.00	12.81	12.82	13.31
$t-stat(Overlap.)$	(2.27)	(2.15)	(2.44)	(1.97)	(1.81)
$t-stat(Nonoverlap.)$	[2.32]	[2.74]	[2.93]	[2.42]	[1.94]
R^2	0.11	0.15	0.20	0.20	0.21
Panel B:					
$A^h dp_t^{Mkt} - dp_t^{Strip}$	13.93	20.87	26.00	26.24	26.05
$t-stat(Overlap.)$	(3.17)	(3.67)	(5.12)	(5.70)	(5.79)
$t-stat(Nonoverlap.)$	[2.67]	[2.67]	[3.79]	[4.49]	[4.90]
R^2	0.12	0.22	0.32	0.34	0.35
N	310	304	298	292	286

Table 5 presents the results of the predictive regressions for the difference between market and strip Sharpe ratios for holding periods ranging from $h = 12$ month through $h = 36$ months. In panel A, the predictor variable is the market dp-ratio. In Panel B, the predictor variable is the scaled difference between the market and strip dp-ratios. We take the average for the strip dp-ratios over the last three months. In parentheses are t-statistics based on the Newey-West (1987) correction with h lags. In brackets are the average t-statistics from h non-overlapping samples starting in months $1, \dots, h$. For each non-overlapping sample, we set the number of lags for the Newey-West correction equal to 3. N reports the number of overlapping observations. The period is from January 1996 through December 2022.

Table 6: **Predicting the Term Structure in Market Model Alphas**

	12m	18m	24m	30m	36m
Panel A:					
dp_t^{Mkt}	-0.01	-0.03	-0.05	-0.06	-0.06
$t-stat(Overlap.)$	(-1.22)	(-1.59)	(-1.82)	(-1.92)	(-1.97)
$t-stat(Nonoverlap.)$	[-0.91]	[-1.52]	[-1.88]	[-2.51]	[-2.85]
R^2	0.06	0.11	0.15	0.17	0.18
Panel B:					
$A^h dp_t^{Mkt} - dp_t^{Strip}$	-0.01	-0.02	-0.04	-0.06	-0.08
$t-stat(Overlap.)$	(-0.80)	(-0.97)	(-1.26)	(-2.09)	(-3.20)
$t-stat(Nonoverlap.)$	[-0.60]	[-0.90]	[-1.12]	[-1.70]	[-2.24]
R^2	0.01	0.02	0.04	0.07	0.11
N	310	304	298	292	286

Table 6 presents the results of the predictive regressions for the term structure of market model alphas (market minus strip alpha) for holding periods ranging from $h = 12$ month through $h = 36$ months. In panel A, the predictor variable is the market dp-ratio. In Panel B, the predictor variable is the scaled difference between the market and strip dp-ratios. We take the average for the strip dp-ratios over the last three months. In parentheses are t-statistics based on the Newey-West (1987) correction with h lags. In brackets are the average t-statistics from h non-overlapping samples starting in months $1, ..., h$. For each non-overlapping sample, we set the number of lags for the Newey-West correction equal to 3. N reports the number of overlapping observations. The period is from January 1996 through December 2022.

Table 7: **Summary of Results**

Main Economic Aspect	Dimension	Finding	Literature (unless new result)
Term Premia	Slope	Rather flat (if anything, insignificantly upward)	BBK: Downward-sloping, insignificant
Term Structure of Sharpe Ratios	Slope	Downward-sloping	BBK: Downward-sloping
Term Structure of Alphas	Slope	Downward-sloping	BBK: Downward-sloping
Term Premia	Cyclicality	Countercyclical	Gormsen (2021): Countercyclical
Term Structure of Sharpe Ratios	Cyclicality	Countercyclical	–
Term Structure of Alphas	Cyclicality	Procyclical	–

Table 7 covers three main economic aspects: (i) term premia (i.e., market minus strip returns), (ii) term structure of Sharpe ratios (i.e., market minus strip Sharpe ratio), and (iii) term structure of alphas (i.e., market minus strip alpha). We assess each economic aspect across the dimensions of slope and its behavior over the business cycle (cyclicality). We record both our empirical findings and older findings in the literature. We compare our empirical results to those reported in BBK and Gormsen (2021), since they use the same definition of term premia.

Figure 1: **Interest Rates**

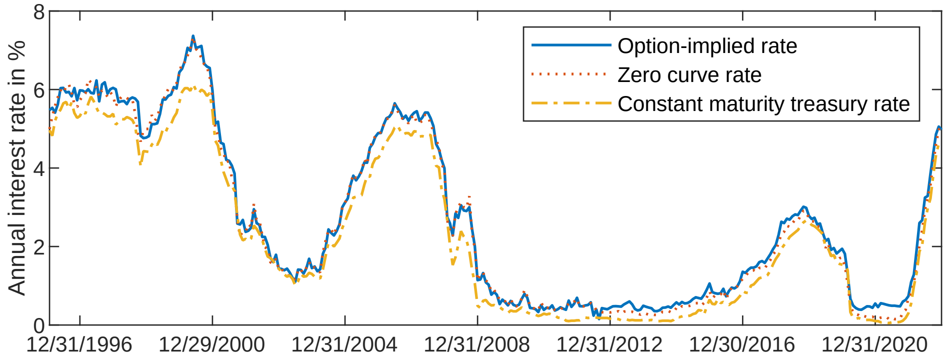

Figure 1 plots the 12-month maturity interest rates. The option-implied rate is based on put-call parity pairs of S&P 500 index options. The zero curve rate is from OptionMetrics. The constant maturity Treasury rate is from the H.15 filing of the St. Louis Federal Reserve Bank. All interest rates are continuously compounded and expressed as a percentage. The period is from January 1996 through December 2022.

Figure 2: **Cumulative Returns**

Panel A: Cumulative returns

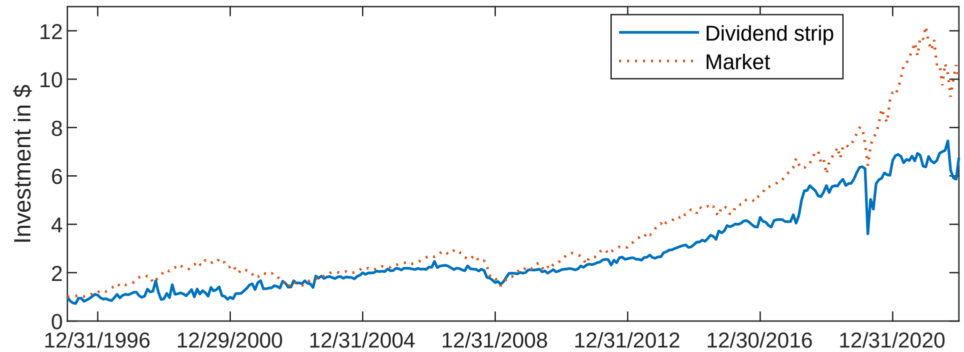

Panel B: Cumulative returns in excess of the risk-free rate

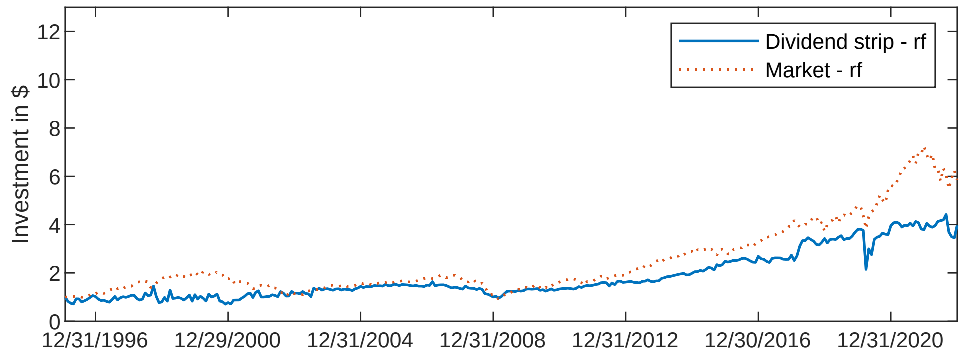

Panel C: Cumulative returns in excess of Treasury returns

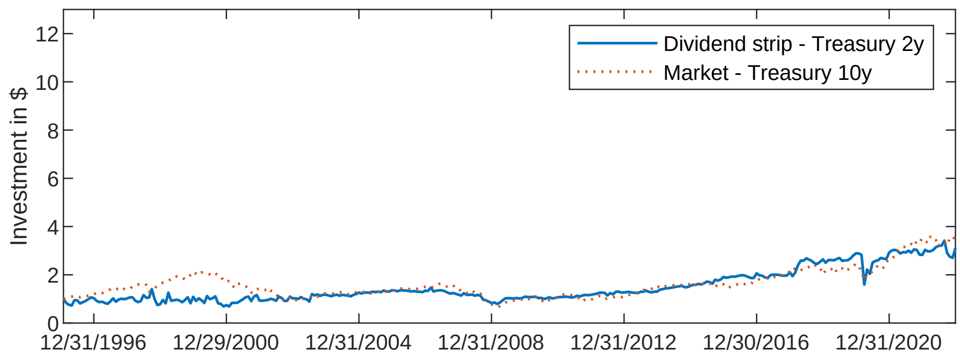

Panel A in Figure 2 plots the cumulative returns for a hypothetical one dollar investment in the dividend strip and the market. Panel B plots the cumulative excess returns of the strip in excess of the risk-free rate and the market in excess of the risk-free rate. Risk-free rate (rf) is the one-month Treasury bill rate. Panel C plots the cumulative excess returns of the strip in excess of the two-year Treasury return and the market in excess of the 10-year Treasury return. The period is from January 1996 through December 2022.

Figure 3: **Annualized Standard Deviation Across Different Holding Periods**

Panel A: Returns in excess of the risk-free rate

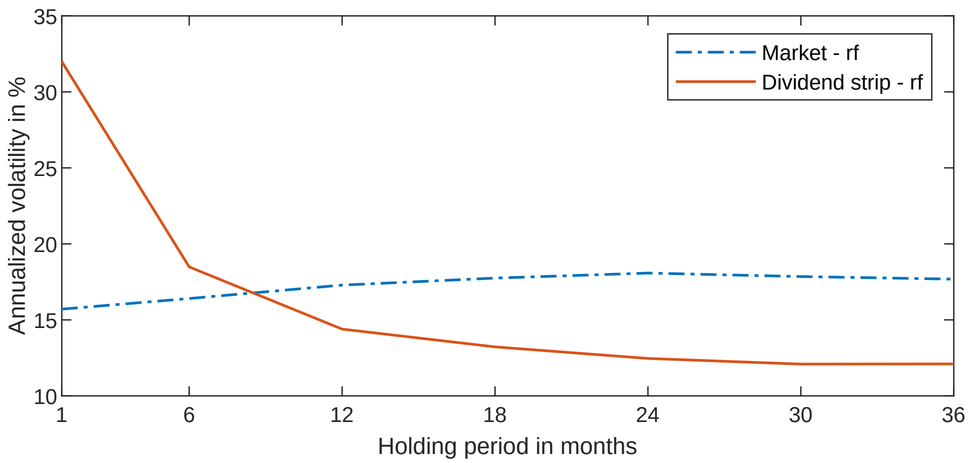

Panel B: Returns in excess of Treasury bond returns

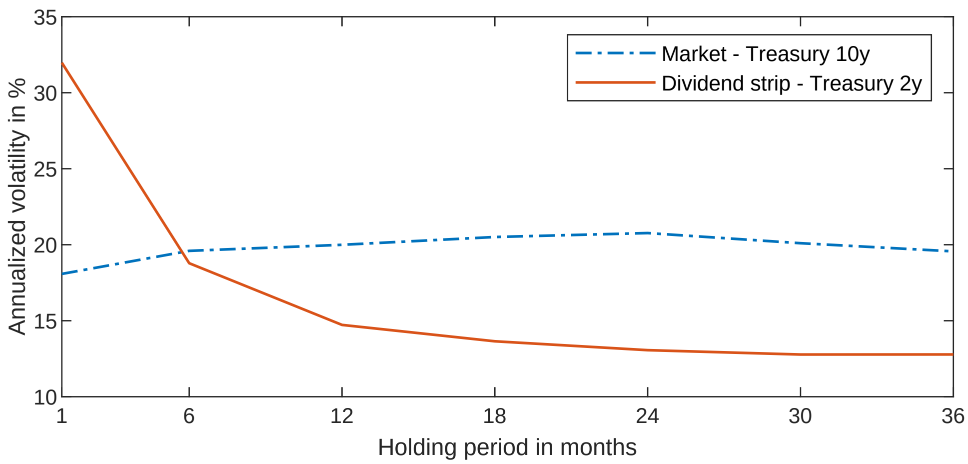

Figure 3 plots the annualized standard deviation for excess market and strip returns for holding periods of 1, 6, 12, 18, 24, 30, and 36 months. The returns are in excess of the risk-free rate (Panel A) and in excess of the Treasury returns (Panel B). The period is from January 1996 through December 2022.

Figure 4: **Annualized Sharpe Ratio Across Different Holding Periods**

Panel A: Returns in excess of the risk-free rate

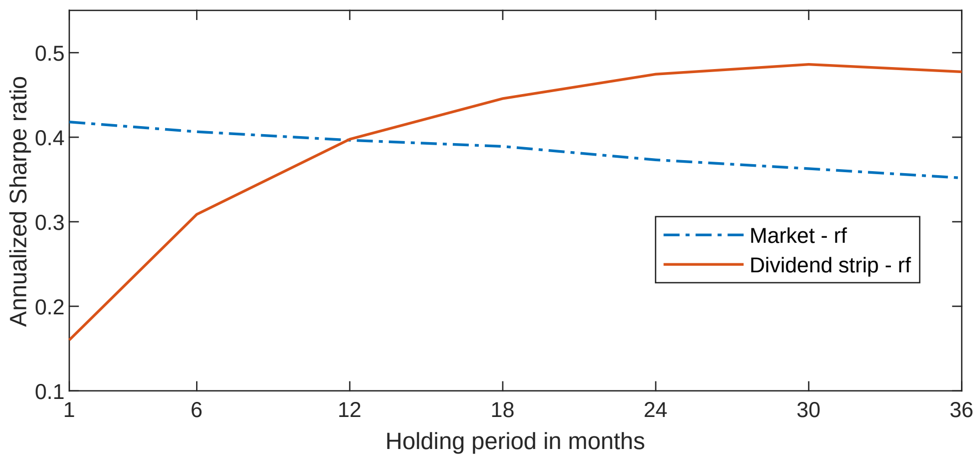

Panel B: Returns in excess of Treasury bond returns

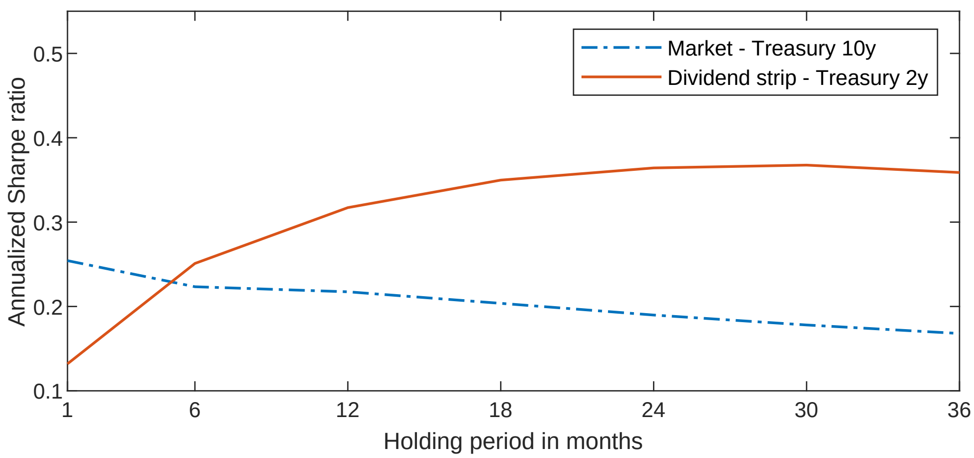

Figure 4 plots the annualized Sharpe ratio for excess market and strip returns for holding periods of 1, 6, 12, 18, 24, 30, and 36 months. The returns are in excess of the risk-free rate (Panel A) and in excess of the Treasury returns (Panel B). The period is from January 1996 through December 2022.

Figure 5: **Annualized Sharpe Ratios: Subsamples**

Panel A: Returns in excess of the risk-free rate

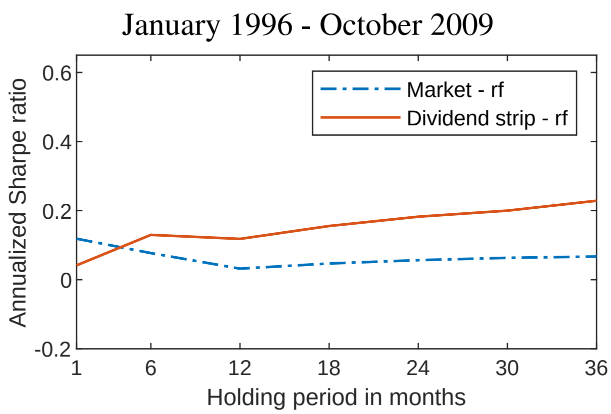

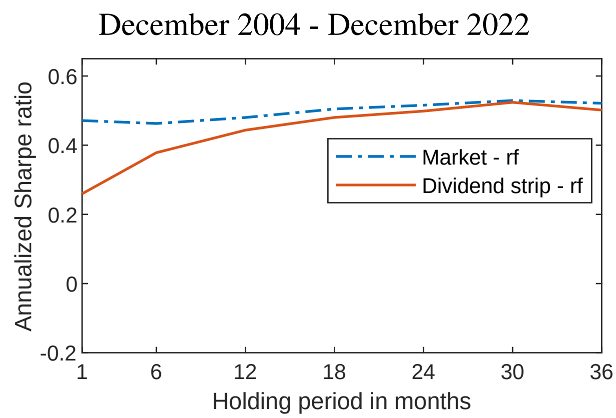

Panel B: Returns in excess of Treasury bond returns

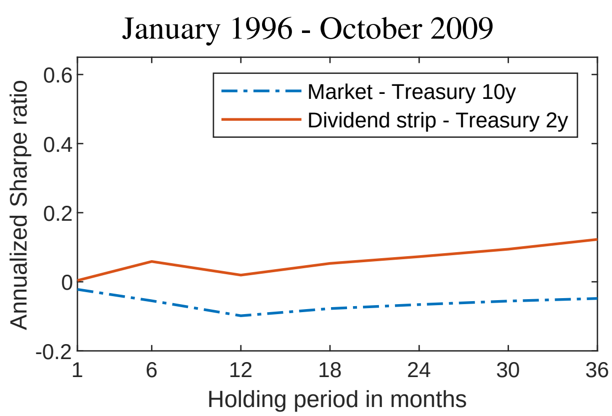

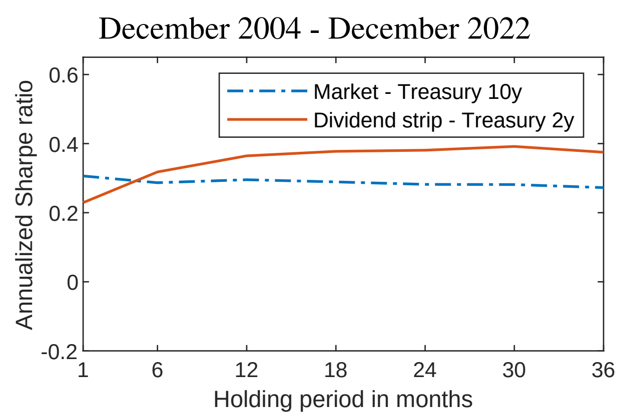

Figure 5 plots the annualized Sharpe ratio for holding periods of 1, 6, 12, 18, 24, 30, and 36 months. The returns are in excess of the risk-free rate (Panel A) and in excess of the Treasury returns (Panel B). The period is from January 1996 through October 2009 (left) and from December 2004 through December 2022 (right).