# Holding Period Effects in Dividend Strip Returns

Benjamin Golez Jens Jackworth *

June 2022

**Abstract**

We estimate short-duration dividend strip prices from 25 years-worth of S&P 500 index option data (1996-2020). We show that short-duration strips offer substantially more attractive returns than does the market, but the measurement error obscures this result at monthly holding periods. For holding periods longer than one year, where the effect of the measurement error dissipates, the strip Sharpe ratio is two to four times the market Sharpe ratio. This outperformance holds in different subperiods, as well as conditionally on recessions or expansions. We also document that the return on the strip in excess of the market is highly predictable.

**JEL Classification: G12, G13, G35**

**Keywords: dividend term structure, dividend strips, option pricing, interest rate invariant**

*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.

# 1 Introduction

Are returns of short-duration equity more attractive than returns of long-duration equity? The answer to this question is important for investment decisions and our understanding of asset pricing models. Yet, the empirical evidence on this topic is mixed. Van Binsbergen, Brandt, and Koijen (2012) estimate short-duration dividend strips from index options during 1996 through 2009 and show that short-duration dividend strips deliver higher returns and Sharpe ratios than does the long-duration market. Bansal, Miller, Song, and Yaron (2021) use dividend futures data from late 2004 to early 2017 and argue the exact opposite. Bansal, Miller, Song, and Yaron (2021) and Gormsen (2021) disagree on whether the relative attractiveness of dividend strips increases or decreases during bad times of the economy. We show that short-duration strip returns are substantially more attractive than are market returns, during both recession and expansions, but performance needs to be measured over long holding periods. Theoretical asset pricing equations typically hold independently of holding periods. However, when prices are measured with error, results can depend on the length of the holding period. This dependence is particularly important for dividend strips estimated from leveraged option positions. Boguth, Carlson, Fisher, and Simutin (2019) show that noise in the option data can cause a nontrivial measurement error in dividend strip returns. They advocate for the use of logarithmic returns to remove the bias in average returns. However, the use of logarithmic returns will not produce unbiased estimates of Sharpe ratios, because the measurement error also inflates the return standard deviation (Blume and Stambaugh 1983). Hence, even with an unbiased estimate for the mean return, the measurement error can still bias Sharpe ratios downwards.

To mitigate the effect of measurement error, we use longer holding period returns. Standard deviation of returns increases over longer holding periods, whereas the effect of measurement error decreases over longer holding periods. Theoretically, the measurement error vanishes asymptotically. Empirically, it disappears for holding periods longer than 12 to 24 months. At these holding periods, we find that dividend strips are more attractive than the market; the dividend strip Sharpe ratio is two to four times the market Sharpe ratio. This difference, which is economically important and statistically significant, is mainly driven by the lower standard deviation of dividend strip returns, whereas average returns on both assets are similar.

In addition to studying returns over longer holding periods, we also lengthen our sample. Bansal, Miller, Song, and Yaron (2021) argue that short samples and an oversampling of recession periods drives many existing results regarding the term structure of equity term premia. We almost double the sample of intradaily option data (1996 through 2020) compared to Van Binsbergen, Brandt, and Koijen (2012) (1996 through 2009). The proportion of recessions in our sample is comparable to the historical occurrence of recessions.[1]

We estimate dividend strip prices from time-matched put-call parity relations (Van Binsbergen, Brandt, and Koijen 2012); (BBK from now on). This estimation is sensitive to the choice of interest rate. The interest rate should reflect the funding costs of marginal investors in the option market (Song 2016; Binsbergen and Koijen 2016). Thus, we suggest an interest rate invariant method and use the endogenous interest rate implied by option prices (Van Binsbergen, Diamond, and Grotteria 2021). This approach contrasts with that of BBK, who impose an exogenous interest rate (zero curve rate), which can bias dividend strip returns. Further, our interest rate invariant approach avoids the time mismatch between the intradaily option data and end-of-day interest rate proxies (Boguth, Carlson, Fisher, and Simutin 2019).

Our main finding is that dividend strip Sharpe ratios for longer holding periods are higher than market Sharpe ratios. This result holds in the full sample, as well as in the subsamples used

[1]Bansal, Miller, Song, and Yaron (2021) find that 15% of their long sample consists of mild recessions and 4% of their sample consists of severe recessions. We have National Bureau of Economic Research (NBER)-defined recessions in 9% of our sample.

by BBK and Bansal, Miller, Song, and Yaron (2021). Moreover, our result holds regardless of whether the investment starts during National Bureau of Economic Research (NBER) recessions or expansions.

Results are robust to the choice of the dividend strip maturity (1.9 years in the base case); the use of the exogenous zero curve rate instead of the interest rate invariant approach; the choice of moneyness as we filter our option sample; the use of raw returns rather than excess returns; and reducing transactions costs by holding the strip to maturity instead of rebalancing monthly.

Having established the attractiveness of dividend strips, we ask whether one can predict the outperformance of the strip over the market. A simple present value model guides our predictive regressions. We find that the difference between the strip and the market return is highly predictable by the difference between the dividend-to-price ratios of the strip and the market. These results strengthen with the length of the holding period and hold both in and out of sample.

Our study bridges the literature on the equity term structure and the one on holding period effects. The equity term structure literature is concerned with the shape of the term structure of returns, standard deviations, and Sharpe ratios (for an overview, see Van Binsbergen and Koijen 2017). Depending on the time period and the measurement of dividend strip prices, researchers argue that dividend strip returns are either more or less attractive than the market (BBK; Bansal, Miller, Song, and Yaron 2021; Gonçalves 2021; Gormsen 2021). We contribute to this literature by constructing a long time series of dividend strip prices. We provide strong evidence that dividend strips deliver economically and statistically higher Sharpe ratios than does the market This is driven by the upward-sloping term structure of standard deviations, while the equity term structure of average returns is mostly flat.

A number of authors challenged the BBK approach of measuring dividend strip prices from the options market. Schulz (2016) suggests that taxes could drive the results, but 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. Further, Binsbergen and Koijen (2016) argue that neither noise induced by asynchronous trading nor alternative funding costs for financial intermediaries are likely explanations for the original findings of BBK. In contrast, we show that measurement error in option prices does affect findings and that results over longer holding periods are not susceptible to such measurement error.

Among the leading asset pricing models, our findings are most consistent with predictions made from the rare disaster model of Gabaix (2009). In this model, short- and long-duration assets are exposed to the same risk and have the same expected returns. However, long-duration returns are more sensitive to the time variation in disaster probabilities. As a result, long-duration assets are more volatile and have lower Sharpe ratios.

The literature on the holding period effects mostly focuses on biases in average returns (Boguth, Carlson, Fisher, and Simutin 2016). We contribute to that literature by focusing on the holding period effects on Sharpe ratios. We show how the measurement error inflates the standard deviation of single-period returns and, thus, requires the use of longer holding period returns to obtain reliable estimates of Sharpe ratios. The measurement error in the monthly returns shows up as a negative AR(1) coefficient and leads to a downward biased Sharpe ratio. This bias disappears over longer holding periods. Our findings dovetail with those of Laarits (2021), who documents the opposite effect for returns on anomalies. Anomalies exhibit positive AR(1) coefficients and upward biased Sharpe ratios. Again, this bias disappears over longer holding periods.

Instead, one could use an ARMA-type model to account for measurement error. Such time-series models suffer when they are misspecified while longer holding period returns work independently of any modeling choice. Thus, longer holding period returns could be an interesting alternative to the main MA(2) model of Getmansky, Lo, and Makarov (2004) that they use to model smoothing of hedge fund returns.

Like BBK, we estimate dividend strip prices from index options. Alternatively, one could deduct dividend strip prices from dividend futures (Van Binsbergen, Hueskes, Koijen, and Vrugt 2013; Cejnek and Randl 2020; Bansal, Miller, Song, and Yaron 2021) Presumably, the use of dividend futures is less prone to measurement error (Binsbergen and Koijen 2016). However, the market for dividend futures has its own set of limitations. It is mostly a sellers' market used for hedging dividend exposure of institutions selling structured products (Dor and Florig 2021). Moreover, exchange-traded dividend futures on the S&P 500 dividend index started quite recently in 2015. Before that, there was only an over-the-counter market for dividend swaps. Most studies using dividend swaps rely on proprietary data from investment banks (e.g., Bansal, Miller, Song, and Yaron 2021;Van Binsbergen, Hueskes, Koijen, and Vrugt 2013) that cannot be freely distributed. In comparison, our estimated dividend strip prices span 25 years of data and can be freely distributed.

Other recent studies estimate the equity term structure from the cross section of equity returns (Giglio, Kelly, and Kozak 2020; Gonçalves 2021). That approach extends the data back even further than by using index options at the cost of estimating the dynamics of the economy and the preferences of the investors. Our data provide direct estimates of dividend strips and can be used as a yardstick to evaluate alternative methods.

Finally, we contribute to the return predictability literature. Researchers have used information on dividends implied by derivative prices (Golez 2014; Li and Wang 2018) or dividend strip returns (BBK) to predict stock market returns. We use present value relations to show that the outperformance of strip returns over market returns is predicted by the scaled difference of their respective dividend-to-price ratios. Recently, Cassella, Golez, Gulen, and Kelly (2021) link the difference in returns to the term structure of investor optimism.

The rest of the paper is organized as follows. In Section 2, we describe the data. In Section 3, we discuss the methodology and estimation of dividend strip prices and returns. In Section 4, we present the results for monthly and longer holding periods. We also discuss the evidence related to subsamples and the business cycle. In Section 5, we predict the term structure of equity returns. In Section 6, we present the results of our robustness checks. Section 7 concludes.

# 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, 2020. 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 2020.

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 2020 from OptionMetrics. Further, we download nominal constant maturity Treasury interest rates from the H.15 filing of the St. Louis Federal Reserve Bank. We also obtain returns on two-year and 10-year fixed maturity Treasuries from CRSP. For recessions, we rely on NBER-defined recessions. 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 short-duration strips compared to the market, we need returns on both securities. We compute market returns from S&P 500 prices. 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 $t$,

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

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

Using exogenous zero curve interest rates, Van Binsbergen, Brandt, and Koijen (2012) invert the put-call parity relation and directly estimate the price $P$ of the short-duration dividend strip from the observed option prices:

$$P_t^\tau = S_t + p_t^\tau(X) - c_t^\tau(X) - Xe^{-rf_t^\tau \tau}. \tag{2}$$

Specifically, for a given day $t$ and maturity $\tau$, 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 an exogenous 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 2019). Second, funding costs of marginal investors in index options may differ from the exogenous interest rate (Song 2016). Ulrich, Florig, and Wuchte (2019) find that the unconditional term premium is either downward-sloping or upward-sloping depending on which interest rate (Treasury-bill rate 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 2). Any mistake in the strip prices is then magnified in the calculation of strip returns. In the Appendix, we consider a simple calibration based on a small error of negative 4 basis points (bp). The error is based on our finding that interest rates implied by the option data (as will be described below) are 4 bp higher than zero curve rates. We show that even such a small error can lead to bias in half-annual dividend strip returns of 0.27% (or 7.63% in relative terms). The elasticity of the strip return with respect to the interest rate error is large at −5.49.

## 3.1 Endogenous Interest Rates

Therefore, we advocate for 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 endogenous variables, the dividend price $P$ and the risk-free rate $rf$. Van Binsbergen, Diamond, and Grotteria (2021) identify the risk-free rate by combining two put-call parity relations with different strike prices $X$ into a pair.[2]

They discuss two different ways of estimating implied interest rates from such pairs. We refer to the first approach as the *outer product approach*. For a given maturity $\tau$, 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 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

$$rf_{t,\tau,a} = -\frac{1}{\tau} ln \left[ \frac{\left(p_t^\tau (X_i) - c_t^\tau (X_i)\right) - \left(p_t^\tau \left(X_j\right) - c_t^\tau \left(X_j\right)\right)}{X_i - X_j} \right]. \tag{3}$$

Finally, we take the median implied rate as the daily 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 $\tau$, we run the following regression based on time-matched put-call parity relations:

$$S_t - c_t^\tau(X) + p_t^\tau(X) = P + \beta X + \epsilon. \tag{4}$$

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 implied interest rates as an input to the put-call parity relation (Equation 2) and 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 sec-

[2] 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.

ond part (2004 through 2020). 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, 2020, the number of unique option relations is 365,009 (84,346 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 implied interest rates from 1996 through 2003 and the regression approach from 2004 through 2020. When we apply both approaches to a sample of recent data, we find that strip prices are virtually the same.[3]

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 5 bp (1.77% 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 33 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 0.79% in relative terms. The difference between both rates increases in the second half of the

[3] 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 slightly noisier dividend prices than the indirect approaches, which we prefer for that reason.

sample to 5.94% in relative terms.[4] This suggests that the zero curve rate was a relatively good proxy for the funding costs of option investors in the early 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 implied interest rates. For comparison, we will report estimations with the zero curve rate in Section 6.

We follow BBK and estimate dividend strip prices using Equation (2). The only difference is that we use 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 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. We estimate dividend strip prices with constant maturities by linearly interpolating the prices for dividend strips just above and below the given maturity.

Finally, we calculate monthly returns on dividend strips. We rely on the approach used by BBK and calculate returns from strip prices with maturities between 1.4 and 1.9 years. We rebalance in January and July. Specifically, each January or July, we buy a dividend strip with a maturity of 1.9 years. Each monthly return is the sum of the strip price at month-end plus the collected dividends, divided by the strip price at the beginning of the month. This is equivalent to reinvesting monthly dividends in dividend strips. We repeat the procedure for 6 months, until we rebalance again into dividend strips with 1.9 years maturity. The only exception to this rule is July 2013 to January 2014, during which we let the strategy rely on maturities between 0.9 and 1.5 years, because the appropriate options maturing in June 2015 were not listed until September 2013.

We calculate monthly returns on the S&P 500 market so that we can compare them with strip

[4] One of the reasons the zero curve rate is lower than the implied interest rate in the recent sample is banks' underreporting of borrowing costs used in the calculation of LIBOR (Gandhi, Golez, Jackwerth, and Plazzi 2019).

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.

## 3.2 Measurement Error

We estimate strip returns from options. Any noise in the option data can lead to biased estimates of performance measures. Boguth, Carlson, Fisher, and Simutin (2019) argue that noise in the data may bias average returns upward and suggest using logarithmic returns. As alluded above, 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, suppose the actual log price $p$ evolves as

$$p_{t+1} = \mu + p_t + \varepsilon_{t+1}, \tag{5}$$

where $\mu$ is the mean log return and $\varepsilon$ is $i.i.d.$ with $\varepsilon \sim N(0, \sigma_\varepsilon^2)$. Thus, the variance of log returns is $\sigma_\varepsilon^2$. Next, suppose that the price is measured with error:

$$\hat{p}_{t+1} = p_{t+1} + \delta_{t+1}, \tag{6}$$

where $\delta$ is $i.i.d.$ with $\delta \sim N(0, \sigma_\delta^2)$. By substitution, we can write the measured price as

$$\hat{p}_{t+1} = \mu + \hat{p}_t + \varepsilon_{t+1} + \delta_{t+1} - \delta_t. \tag{7}$$

Assuming that returns and measurement errors are uncorrelated (i.e., $cov(\varepsilon, \delta) = 0$), the AR(1) coefficient $\Phi$ is

$$\Phi = \frac{cov(\hat{r_{t+1,}}, \hat{r}_t)}{var(\hat{r}_t)} = \frac{-\sigma_\delta^2}{2\sigma_\delta^2 + \sigma_\varepsilon^2}. \tag{8}$$

Thus, measurement error leads to a negative serial correlation in returns. This negative serial correlation inflates the variance of measured returns.

To reduce the influence of measurement error, we focus on longer holding periods. Define $h$-period returns as $r_{t,t+h} = \sum_{j=1}^{h} r_{t+j}$. Variance of actual returns scales by the time horizon, whereas variance due to measurement error stays constant across the holding periods. As a result, the variance bias decreases in the length of the holding period. For single period returns, the ratio of measured variance to actual variance is:

$$\frac{var(\hat{r_{t,t+1}})}{var(r_{t,t+1})} = \frac{2\sigma_\delta^2 + \sigma_\varepsilon^2}{\sigma_\varepsilon^2} = 1 + 2\frac{\sigma_\delta^2}{\sigma_\varepsilon^2}. \tag{9}$$

For $h$-period returns, the same ratio is:

$$\frac{var(\hat{r_{t,t+h}})}{var(r_{t,t+h})} = 1 + \left(\frac{2}{h}\right)\frac{\sigma_\delta^2}{\sigma_\varepsilon^2}. \tag{10}$$

Hence, in the limit, an infinitely long holding period would completely eliminate the variance bias due to measurement error. We can express the ratio of h-period measured Sharpe ratio to actual Sharpe ratio as:

$$\frac{SR(\hat{r}_t)}{SR(r_t)} = \frac{E[r_{t,\hat{t}+h}] - rf_{t,t+h}}{\sqrt{(2\sigma_\delta^2 + h\sigma_\varepsilon^2)}} \frac{\sqrt{(h\sigma_\varepsilon^2)}}{E[r_{t,\hat{t}+h}] - rf_{t,t+h}} = \frac{\sqrt{(h\sigma_\varepsilon^2)}}{\sqrt{(2\sigma_\delta^2 + h\sigma_\varepsilon^2)}} = \frac{1}{\sqrt{1 + \left(\frac{2}{h}\right)\frac{\sigma_\delta^2}{\sigma_\varepsilon^2}}} = \frac{1}{\sqrt{\frac{var(r_{t,\hat{t}+h})}{var(r_{t,t+h})}}}. \tag{11}$$

Thus, the ratio of Sharpe ratios approaches one from below as $h$ grows. For sufficiently long holding period, the measured Sharpe ratio will approach the actual Sharpe ratio. In the empirical analysis, we will consider holding periods of up to 60 months.

### 3.3 Dividend-to-Price Ratios

We define the log dividend-to-price ratio for the market $dp^{Mkt}$ as the logarithm of the current level of the S&P 500 index minus the logarithm of the sum of dividends over the past year. For dividend strips, we define $dp^{Strip}$ as the logarithm of the price of a one-year dividend strip minus the logarithm of the sum of dividends over the past year.

## 4 Results

Table 1 reports the summary statistics for single-period (monthly) returns. Columns 1 and 3 give statistics for dividend strip returns and market returns. In columns 2 and 4, we compute the equity risk premium by subtracting the returns on fixed maturity Treasuries from equity returns. We want to approximately match the duration of equities and Treasuries. For dividend strips, we use two-year maturity notes. For the market, we follow Van Binsbergen and Koijen (2017) and use 10-year bonds.[5] All returns in Table 1 are in logarithms. Figure 2 plots the cumulative returns for rolling over investments in the dividend strip or the market.

[5]Results are qualitatively similar if we use 20-year bonds.

[Table 1 about here]

[Figure 2 about here]

We only find small differences between the strip and the market in terms of average returns. The average annualized strip return is 7.63%, whereas the average market return is 9.02%.[6] As depicted in Panel A in Figure 2, $1.00 invested in the dividend strip for 25 years grows to $6.68, whereas $1.00 invested in the market grows to $9.46.

When we subtract the Treasury returns, the average strip return in excess of the two-year Treasury return is 4.34% annually and the market return in excess of the 10-year Treasury return is 3.91% annually. Panel B in Figure 2 shows that $1.00 invested in the strip in excess of the two-year Treasury would grow to $2.96, and the $1.00 invested in the market in excess of the 10-year Treasury would grow to $2.77. Thus, accounting for the Treasury term structure matters for the relative comparison of strips and the market, but neither the difference between the average returns on the strip and the market nor the difference between average excess returns is statistically significant.

While we do not find much difference in terms of average returns, short-duration dividend strip returns are substantially more volatile than market returns. Regardless of whether or not we deduct the Treasury returns, dividend strip returns are approximately twice as volatile as market returns. The standard deviation of single-period dividend strip returns is 33%, whereas the standard deviation of single-period market returns is 19% (16% if we do not subtract Treasury returns). Monthly dividend strip returns also exhibit a strong negative serial autocorrelation of -0.34. This means that lagged strip returns explain 12% of the variation in monthly returns. In comparison,

[6] These results are contrary to those of Van Binsbergen, Brandt, and Koijen (2012), who find that dividend strips offer higher returns than the market. In the Internet Appendix, we replicate their results and compare them to our estimates. We find that the difference stems from both our use of the interest rate invariant approach and our extended sample. With the interest rate invariant approach, dividend strips and the market deliver similar returns during the BBK sample, but dividend strips underperform during the extended sample (1996 through 2020).

the AR(1) coefficient for the market return is almost zero and the AR(1) coefficient for the market in excess of Treasury return is slightly positive at 0.06.[7]

## 4.1 Longer Holding Periods

A high standard deviation of single-period strip returns combined with a strong negative serial correlation in returns is indicative of a measurement error in dividend strip prices. In Section 3.2, we show that the effect of the measurement error declines as the holding period over which we calculate performance measures increases.

Next we consider holding periods of 1 through 60 months. We sum the logarithmic returns over a given holding period, that is, $r_t^h = \sum_{j=1}^{h} r_{t+1-j}$, for $h = 1, \dots, 60$. We focus on excess returns, that is, dividend strip returns minus 2-year Treasury returns and market returns minus 10-year Treasury returns. Figure 3 presents annualized standard deviations across different holding periods. Table 2 reports the corresponding summary statistics.

[Table 2 about here]

[Figure 3 about here]

We note a drastic decrease in the annualized standard deviation for the excess dividend strip. The standard deviation decreases from 33% for monthly returns to 15% for annual returns. It then stabilizes at around 13% for holding periods beyond two years. This suggests that obtaining stable estimates for the standard deviation of dividend strip returns takes at least 12 to 24 periods. In comparison, the standard deviation for the market in excess of the 10-year bond initially increases slightly from 18% to 20% and then decreases to 17%. Overall, these patterns are consistent with a

[7] When we check for a higher-order serial correlation, 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.

strong negative serial correlation for the dividend strip and a slightly positive serial correlation for the market.

[Figure 4 about here]

These shifts in the standard deviation profoundly affect the annualized Sharpe ratios (see Figure 4). Specifically, the dividend strip Sharpe ratio increases from 0.13 for the monthly holding period to 0.35 for the five-year holding period. In contrast, the market Sharpe ratio decreases from 0.21 to around 0.10. We test for the difference in Sharpe ratios using the heteroskedasticity- and autocorrelation-consistent (HAC) test proposed by Ledoit and Wolf (2008). We find that the strip Sharpe ratio is significantly higher than the market Sharpe ratio for any holding period longer than 24 months. Since these are the holding periods in which the measurement error is minimized, we conclude that the dividend strip offers more attractive returns than does the market.[8,9]

## 4.2 Subsamples and the Business Cycle

How does the dividend strip return compare with the market return over different subsamples? Panel A of Figure 5 plots the annualized Sharpe ratios during the BBK period from January 1996 through October 2009. Panel B plots the results for the period from December 2004 through December 2020 (in Bansal, Miller, Song, and Yaron (2021), the time period is from December

[8] In the Internet Appendix, we show that Sharpe ratios exhibit similar pattern if we use raw market and dividend strip returns rather than returns in excess of bond returns (A.2).

[9] One observation requires additional consideration. For the market, the decrease in the Sharpe ratio goes beyond the increase in its standard deviation. In fact, the decrease in the Sharpe ratio for the market is mostly driven by the decrease in the average market return. This is because, for longer holding periods, when using overlapping observations, we place more (less) weight on the observations from the center (early and late) years of the sample period. As a robustness check, we consider a circular bootstrap as in Politis and Romano (1992), where we connect the last return with the first. While unrealistic from the point of view of an investor, this approach ensures that the average returns are the same regardless of the holding period. We still find that the difference in Sharpe ratios of excess strip and market returns is statistically significant for longer holding periods.

2004 to February 2017) In both subsamples, the dividend strip Sharpe ratio is higher than the market Sharpe ratio for all holding periods and substantially so at longer holding periods.

[Figure 5 about here]

[Figure 6 about here]

How does the dividend strip return compare with the market return over business cycle? We use standard NBER business cycle dating. If the first return occurs during an expansion, we compute holding period returns out to 60 months and report their Sharpe ratios in Panel A of Figure 6. The pattern closely follows the unconditional results of Figure 4, in particular for longer holding periods. For shorter holding periods of up to about half a year, the market outperforms the strip. Thereafter, as in the unconditional results, the strip outperforms the market in terms of Sharpe ratios. Panel B of Figure 6 shows the Sharpe ratios when the first holding period return occurs during a recession. During recessions, the strip outperforms the market for all holding periods. Again, for long holding periods, the Sharpe ratios converge to the unconditional results of Figure 4.

## 4.3 Discussion of Results

Our results show that the strip outperforms the market in terms of Sharpe ratios, as long as we focus on longer holding periods, where the measurement error is minimized. What are the economic drivers of this outperformance?

The fact that Sharpe ratios are higher for short-duration dividend strips than for the market implies a downward-sloping term structure of Sharpe ratios. BBK study term structures of Sharpe ratios for a set of popular asset pricing models, including the Bansal and Yaron (2004) long-run risks model, the Campbell and Cochrane (1999) habit formation model, the Gabaix (2009) rare disaster model, and the Lettau and Wachter (2007) value premium model. They depict simulated Sharpe ratios in their figures 5 and 6 (both panel C). As the theoretical models exhibit no measurement error, their Sharpe ratio term structures hold for all holding periods. The long-run risks and the habit formation models both exhibit an upward-sloping term structure of Sharpe ratios and are thus inconsistent with our findings. The Gabaix (2009) rare disaster model and the Lettau and Wachter (2007) value premium model both show downward-sloping term structures of Sharpe ratios consistent with our empirical findings.

Of the latter two models, the rare disaster model of Gabaix (2009) seems to fit our findings best. In the model, short- and long-duration assets are exposed to the same risk of a rare disaster and, hence, have the same expected return. Meanwhile, the return standard deviation is increasing with maturity since long-maturity assets are more exposed to the time variation in disaster probabilities. As a result, the Sharpe ratio is downward-sloping. This is exactly what we find in our analysis. The difference in average returns between short-duration dividend strips and long-duration market is small. The annualized standard deviation of market returns, however, is higher than the annualized standard deviation for the strips (i.e., 20% vs. 13% at the 24-month holding period). Thus, the rare disaster model seems to capture our results not only in terms of Sharpe ratios (the focus of our study) but also in terms of average returns and standard deviations.

The value premium model of Lettau and Wachter (2007) captures our results in terms of the Sharpe ratios, but it predicts a downward-sloping term structure of both expected returns and standard deviations, whereas we find that the standard deviation of the long-duration market is higher than the standard deviation of the short-duration dividend strip.

# 5 Predictability of the Outperformance of the Strip Over the Market

So far, we have shown that the strip outperforms the market in terms of the Sharpe ratio. In this section, we investigate whether the difference between returns on the strip and the market is predictable. We motivate our analysis with a simple 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, which we assume to be an AR(1) process:

$$E\left(\Delta d_{t+1}\right) = g_t = \gamma_0 + \gamma_1 g_{t-1} + \varepsilon_t^g. \tag{12}$$

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 follow an AR(1) process:

$$E\left(r_{t+1}^{Mkt}\right) = \mu_t^{Mkt} = \delta_0 + \delta_1 \mu_{t-1}^{Mkt} + \varepsilon_t^{\mu,Mkt}. \tag{13}$$

For expected strip returns, we do not impose any specific dynamics:

$$E\left(r_{t+1}^{Strip}\right) = \mu_t^{Strip}. \tag{14}$$

Under these assumptions and using the Campbell and Shiller (1988) decomposition, we can write the logarithm of the market dividend-to-price ratio as (Van Binsbergen and Koijen 2010)

$$dp_t^{Mkt} = \kappa + \left(\frac{1}{1-\rho\delta_1}\right)\mu_t^{Mkt} - \left(\frac{1}{1-\rho\gamma_1}\right)g_t, \tag{15}$$

where $\rho = \frac{exp(-\bar{dp})}{1+exp(-\bar{dp})}$. The logarithm of the strip dividend-to-price ratio is:

$$dp_t^{Strip} = \mu_t^{Strip} - g_t. \tag{16}$$

Then the outperformance is the difference between expected strip and market returns:

$$E\left(r_{t+1}^{Strip}\right) - E\left(r_{t+1}^{Mkt}\right) = A\kappa + \left[dp_t^{Strip} - Adp_t^{Mkt}\right] + Bg_t, \tag{17}$$

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

## 5.1 In-Sample Predictability

We use Equation (17) to motivate our predictive regressions. We start with univariate regressions of differences in strip and market returns over the next $h$ periods on either the market or the strip dividend-to-price ratio:

$$\sum_{j=1}^{h}\left(r_{t+j}^{Strip} - r_{t+j}^{Mkt}\right) = \alpha + \beta X_t + \varepsilon, \tag{18}$$

where $h = 1,..,60$ months, and $X_t$ is either $dp_t^{Strip}$ or $dp_t^{Mkt}$.

[Table 3 about here]

Panels A and B of Table 3 report the results. Newey-West t-statistics appear in parentheses, with the number of lags equal to $h$. In brackets are t-statistics based on $h$ nonoverlapping observations. Specifically, for a given holding period, we estimate the predictive regression on $h$ alternative nonoverlapping samples. We then average $t$-statistics across the $h$ alternative samples and report them in the table. The strip dividend-to-price ratio predicts the realized term structure over shorter periods (coefficients are significant for holding periods shorter than three years), whereas the market dividend-to-price ratio predicts the realized term structure over longer holding periods (coefficients are statistically significant for holding periods longer than two years).

Next, we consider a univariate regression of the realized term structure on a scaled difference between the dividend-to-price ratios as implied by Equation (17):

$$\sum_{j=1}^{h}\left(r_{t+j}^{Strip}-r_{t+j}^{Mkt}\right)=\alpha+\beta_{SD}\left[dp_t^{Strip}-\left(1-\rho\delta_1^h\right)dp_t^{Mkt}\right]+\upsilon. \tag{19}$$

The scaling factor depends on $\rho$, $\delta_1$, and $h$. During our sample period, the coefficient $\rho$ is 0.9825. To estimate the persistence of expected returns, we follow Golez and Koudijs (2020) and infer the persistence of expected returns $\beta(h)$ from a predictive regression of market returns over the next $h$ months on lagged values of the market dividend-to-price ratio as $\beta(h)=\frac{1-\delta_1^{h/12}}{1-\delta_1}\beta(12)$. We estimate $\hat{\beta}(60)=1.16$ and $\hat{\beta}(12)=0.34$, which implies a persistence of expected returns of $\delta_1=0.81$

Panel C of Table 3 reports the results. Unlike the univariate dividend-to-price ratio regressions that predict the term structure either over the short horizon ($dp_t^{Strip}$) or the long horizon ($dp_t^{Mkt}$), the scaled difference between the two dividend-to-price ratios predicts the equity term structure at any holding period. The $R$-squared varies from 11% at the monthly holding period to 25% at the annual holding period to 66% at the five-year holding period.

Finally, in the last regression specification, guided by Equation (17), we add a proxy for expected dividend growth:

$$\sum_{j=1}^{h}\left(r_{t+j}^{Strip}-r_{t+j}^{Mkt}\right)=\alpha+\beta_{SD}\left[dp_t^{Strip}-\left(1-\rho\delta_1^h\right)dp_t^{Mkt}\right]+\beta_g g_t+\eta. \tag{20}$$

We proxy for expected dividend growth by the logarithm of indicated dividends over the 12-month trailing sum of dividends. The indicated dividends over the next year are provided by the S&P Dow Jones Indices. They are based on announced dividends, or, if dividends have not yet been announced, they are the last announced dividends projected into the future.

Panel D in Table 3 reports the results. When we add the indicated dividend growth as an additional predictor, the $R$-squared increases compared to Panel C to 16% at the monthly holding period to 32% at the annual holding period to 72% at the five-year holding period. The estimated coefficient for the indicated dividend growth is positive, which is consistent with the widely documented feature that expected returns are more persistent than expected growth rates (Golez and Koudijs 2020). Specifically, in Equation (17), the coefficient $B$ is positive iff $\delta_1 > \gamma_1$.

## 5.2 Out-of-Sample Predictability

While in-sample predictability is strong, particularly at longer holding periods, in-sample results do not necessarily imply that returns are predictable in real time (Goyal and Welch 2003; Goyal and Welch 2008; Cochrane 2008). We check this by computing the out-of-sample R-square (ROOS) as in Goyal and Welch (2008):

$$ROOS=1-\frac{\sum_{\tau=1}^{T}\left(\left(r_\tau^{Strip}-r_\tau^{Mkt}\right)-\left(r_\tau^{\hat{Strip}}-r_\tau^{\hat{Mkt}}\right)\right)^2}{\sum_{\tau=1}^{T}\left(\left(r_\tau^{Strip}-r_\tau^{Mkt}\right)-\left(\overline{r_\tau^{Strip}}-\overline{r_\tau^{Mkt}}\right)\right)^2}, \tag{21}$$

where $r^{Strip}-r^{Mkt}$is the difference between the actual return on the strip and the market, $r^{\hat{Strip}}-r^{\hat{Mkt}}$ is the difference between the predicted return on the strip and the market estimated on the sample up to $\tau - 1$, and $\overline{r^{Strip}} - \overline{r^{Mkt}}$ is the mean return up to $\tau - 1$. We use 60 months for the first training sample and Eq (21) with the scaled difference between the dividend-to-price ratio for the strip and the market (as in Eq (19)) to make out-of-sample predictions. Our training sample limits the lengths of holding periods that we can reasonably estimate. We limit ourselves to 12- and 24-month ahead returns, for which we compute ROOS and Clark and West (2007) t-statistics. For 12-months holding period returns, the ROOS is 15.58% and significant with a *t*-statistics of 2.65. Results are similar for 24-months holding period returns with an ROOS of 14.49% and a *t*-statistics of 2.31. These results are much stronger than typical ROOS values in the literature and suggest that the predictability can be exploited in real time.

# 6 Robustness

We consider several robustness checks with respect to dividend strip maturity, exogenous interest rates, option moneyness, and transaction costs. Table 4 repeats the main strip result (Table 2, Panel A) and collects the robustness results.

[Table 4 about here]

## 6.1 Dividend Strip Maturity

In the base case, we invest each January in a dividend strip with an approximate maturity of 1.9 years. We collect dividends each month and hold this position for half a year until we rebalance into a new 1.9-year dividend strip in July. We now check our results for different maturities of the strip. Panels B and C of Table 4 show results for maturities of 1.3 and 0.9 years; the results are very similar to the base results in Panel A for a maturity of 1.9 years, while the average returns and standard deviations are slightly higher than in the base case, Sharpe ratios are very close to the base case Sharpe ratios and are always increasing in the length of the holding period.

## 6.2 Exogenous Interest Rates

In the base case, we estimate dividend strip prices using an interest rate invariant approach. Now, we use zero curve interest rates instead. In Panel D of Table 4, we find that excess returns are about 1% higher than in the base case, leading to a relative error of 37%. Standard deviations are somewhat higher than in the base case (except for monthly returns, where the standard deviation is lower). As a result, Sharpe ratios are overestimated with a relative error of about 30% (even more so for monthly Sharpe ratios). However, regardless of the bias in the level of Sharpe ratios, we observe the same pattern of increasing Sharpe ratios in the length of the holding period from 0.21 at the monthly holding period to 0.44 at the five-year holding period.

The error in the level of Sharpe ratios aligns with our earlier analysis that the zero curve rate underestimates the interest rate of marginal investors in the option markets and, thus, underestimates the prices of dividend strips, which leads to an upward bias in dividend strip returns. In the appendix, we show the mechanics of the upward bias and that using constant maturity Treasury rates further amplifies errors in strip returns.[10] We interpret these results as an additional argument for the use of the interest rate invariant approach.

## 6.3 Option Moneyness

We identify implied interest rates by combining put-call pairs with different strike prices. In the base case, we use a wide range of strike prices to estimate implied rates with moneyness levels

[10]We do not use overnight interest rates or repo rates, as neither matches the maturity range of our dividend strip (from 1.4 through 1.9 years) and are only available in recent years.

between 0.5 and 1.5. We use the same range of moneyness levels in the calculation of dividend strips. We now consider the same time series except that, for the period from 2004 onward, we use only options with moneyness levels between 0.8 and 1.2. Panel E of Table 4 shows that Sharpe ratios are increasing in the length of the holding period, and all the main conclusions remain the same. For longer holding periods, returns are a little higher and have a somewhat lower standard deviation than in the base case. As a result, Sharpe ratios are somewhat higher at longer holding periods than in the base case in Panel A.

### 6.4 Transaction Costs

Our trading strategy consists of buying a 1.9-year dividend strip at the end of January. We collect dividends over the month of February, sell the asset at the end of February, and compute our February return. Then we buy back the asset (or never sell it), collect dividends over the next month, and sell the asset again. We repeat this strategy for 6 months until the end of July, when the strip has a maturity of 1.4 years. Thereafter, we rebalance into a new 1.9-year strip. Based on this time series of monthly returns, we calculate cumulative returns (see Figure 2). Although we only rebalance into a new asset every 6 months, the cumulative strategy assumes that dividends are reinvested in the dividend strip every month. As the strategy involves options with large bid/ask spreads, we are concerned about trading costs. We follow Bansal, Miller, Song, and Yaron (2021) and consider holding the strip to maturity (instead of rebalancing monthly) to reduce trading costs.

More precisely, every January and July, we buy a dividend strip with a maturity of 1.9 years and hold it until maturity. We collect monthly dividends and reinvest them in the S&P 500 index. Our return is the logarithm of the value of reinvested dividends over the initial price. For comparison, we similarly calculate returns on the S&P 500 index as the logarithm of the future S&P 500 price, plus the value of reinvested dividends over the initial S&P 500 price.

For a better comparison with the results in Section 4, we present excess returns. The holding period of the hold-to-maturity strategy is 1.9 years (22.8 months) and is thus comparable to the main results for a holding period of 24 months in Table 2. The average annualized excess returns for the strip (4.31%) and the market (2.47%) are slightly lower than in the base case. The standard deviations for the strip excess returns (14.29%) and the market (20.65%) are slightly higher than in the base case. The resultant Sharpe ratios are 0.30 and 0.12 and slightly lower than in the base case. The strip Sharpe ratio is thus 2.6 times higher than the market Sharpe ratio. This relative outperformance of the strip over the market in terms of the Sharpe ratio is very similar to our main analysis. Based on the arguments of Bansal, Miller, Song, and Yaron (2021), we conclude that our results are robust to the presence of transaction costs.

# 7 Concluding Remarks

We estimate dividend strip prices from intradaily data for options on the S&P 500 index from 1996 to 2020. We almost double the existing time series of strip prices, suggest an interest rate invariant approach 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. While researchers have previously argued that Sharpe ratios for short-duration assets are either higher or lower than market Sharpe ratios, we show that the strip clearly dominates the market, as long as we focus on longer holding periods, where the effect of the measurement error is marginal. For holding periods longer than a year, the standard deviation for strip returns is smaller than the market standard deviation, and the strip Sharpe ratio significantly exceeds the market Sharpe ratio. Such outperformance of strips holds during the full sample, in subsamples, and during both expansions and recessions. Overall, our results are most consistent with the predictions of the rare disaster model of Gabaix (2009), which predicts a flat term structure of expected returns, an upward-sloping term structure for standard deviations, and a downward-sloping term structure for Sharpe ratios. In additional results, we show that the outperformance of the strip over the market is highly predictable by the relative prices of short- and long-duration assets. This predictability is strongest at longer holding periods and holds both in- and out-of-sample.

# References

Bansal, Ravi, Shane Miller, Dongho Song, and Amir Yaron (2021). "The term structure of equity risk premia." *Journal of Financial Economics* 142 (3), 1209–1228.

Bansal, Ravi and Amir Yaron (2004). "Risks for the long run: A potential resolution of asset pricing puzzles." *Journal of Finance* 59 (4), 1481–1509.

Binsbergen, Jules H van and Ralph SJ Koijen (2016). "On the timing and pricing of dividends: reply." *American Economic Review* 106 (10), 3224–37.

Blume, Marshall E and Robert F Stambaugh (1983). "Biases in computed returns: An application to the size effect." *Journal of Financial Economics* 12 (3), 387–404.

Boguth, Oliver, Murray Carlson, Adlai Fisher, and Mikhail Simutin (2016). "Horizon effects in average returns: The role of slow information diffusion." *Review of Financial Studies* 29 (8), 2241–2281.

Boguth, Oliver, Murray Carlson, Adlai J Fisher, and Mikhail Simutin (2019). "Levered noise and the limits of arbitrage pricing: Implications for the term structure of equity risk premia." *Working paper, Arizona State University*.

Campbell, John Y and John H Cochrane (1999). "By force of habit: A consumption-based explanation of aggregate stock market behavior." *Journal of Political Economy* 107 (2), 205–251.

Campbell, John Y and Robert J Shiller (1988). "The dividend-price ratio and expectations of future dividends and discount factors." *Review of Financial Studies* 1 (3), 195–228.

Cassella, Stefano, Benjamin Golez, Huseyin Gulen, and Peter Kelly (2021). "Horizon bias and the term structure of equity returns." *Working paper, Tilburg University*.

Cejnek, Georg and Otto Randl (2020). "Dividend risk premia." *Journal of Financial and Quantitative Analysis* 55 (4), 1199–1242.

Clark, Todd E and Kenneth D West (2007). "Approximately normal tests for equal predictive accuracy in nested models." *Journal of Econometrics* 138 (1), 291–311.

Cochrane, John H (2008). "The dog that did not bark: A defense of return predictability." *Review of Financial Studies* 21 (4), 1533–1575.

Demaskey, Andrea L and Jean L Heck (1998). "Put-call-forward exchange parity and the implied riskfree rate." *Advances in Investment Analysis and Portfolio Management Vol 5, ed. Lee, C-F, Oxford, UK: JAI Press*, 70–85.

Dor, Arik Ben and Stephan Florig (2021). "Euro Stoxx 50 DividendsâReconciling Analyst Estimates and Dividend Future Prices." *Journal of Portfolio Management* 47 (7), 63–77.

Gabaix, Xavier (2009). "Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance." *Working paper, New York University*.

Gandhi, Priyank, Benjamin Golez, Jens Carsten Jackwerth, and Alberto Plazzi (2019). "Financial market misconduct and public enforcement: The case of Libor manipulation." *Management Science* 65 (11), 5268–5289.

Getmansky, Mila, Andrew W Lo, and Igor Makarov (2004). "An econometric model of serial correlation and illiquidity in hedge fund returns." *Journal of Financial Economics* 74 (3), 529–609.

Giglio, Stefano, Bryan T Kelly, and Serhiy Kozak (2020). "Equity term structures without dividend strips data." *Working paper, Yale University*.

Golez, Benjamin (2014). "Expected returns and dividend growth rates implied by derivative markets." *Review of Financial Studies* 27 (3), 790–822.

Golez, Benjamin and Peter Koudijs (2020). "Equity duration and predictability." *Working paper, University of Notre Dame*.

Gonçalves, Andrei (2021). "Reinvestment risk and the equity term structure." *Journal of Finance* 76 (5), 2153–2197.

Gormsen, Niels J (2021). "Time variation of the equity term structure." *Journal of Finance* 76 (4), 1959–1999.

Goyal, Amit and Ivo Welch (2003). "Predicting the equity premium with dividend ratios." *Management Science* 49 (5), 639–654.

— (2008). "A comprehensive look at the empirical performance of equity premium prediction." *Review of Financial Studies* 21 (4), 1455–1508.

Laarits, Toomas (2021). "Discounting market timing strategies." *Working Paper, New York University*.

Ledoit, Oliver and Michael Wolf (2008). "Robust performance hypothesis testing with the Sharpe ratio." *Journal of Empirical Finance* 15 (5), 850–859.

Lettau, Martin and Jessica A Wachter (2007). "Why is long-horizon equity less risky? A duration-based explanation of the value premium." *Journal of Finance* 62 (1), 55–92.

Li, Ye and Chen Wang (2018). "Rediscover predictability: Information from the relative prices of long-term and short-term dividends." *Fisher College of Business Working Paper* (2018-03), 016.

Politis, Dimitris N and Joseph P Romano (1992). "A circular block-resampling procedure for stationary data." *Exploring the Limits of Bootstrap, eds. R LePage and L Billard, New York: John Wiley*, 263–270.

Schulz, Florian (2016). "On the timing and pricing of dividends: Comment." *American Economic Review* 106 (10), 3185–3223.

Sialm, Clemens (2009). "Tax changes and asset pricing." *American Economic Review* 99 (4), 1356–83.

Song, Yang (2016). "Dealer funding costs: Implications for the term structure of dividend risk premia." *Working paper, University of Washington*.

Ulrich, Maxim, Stephan Florig, and Christian Wuchte (2019). "A model-free term structure of US dividend premiums." *Working paper, Karlsruhe Institute of Technology*.

Van Binsbergen, Jules, Michael Brandt, and Ralph Koijen (2012). "On the timing and pricing of dividends." *American Economic Review* 102 (4), 1596–1618.

Van Binsbergen, Jules, William F Diamond, and Marco Grotteria (2021). "Risk-free interest rates." *Journal of Financial Economics, forthcoming*.

Van Binsbergen, Jules, Wouter Hueskes, Ralph Koijen, and Evert Vrugt (2013). "Equity yields." *Journal of Financial Economics* 110 (3), 503–519.

Van Binsbergen, Jules and Ralph Koijen (2010). "Predictive regressions: A present-value approach." *Journal of Finance* 65 (4), 1439–1471.

— (2017). "The term structure of returns: Facts and theory." *Journal of Financial Economics* 124 (1), 1–21.

# Appendix

In this appendix, we provide a simple calibration exercise to illustrate the effect of errors in the risk-free rate on dividend strip returns according to Equation (2). We thank an anonymous referee for the suggestion to use the observed term structure at 1.4 and 1.9 years. In line with Figure 1, both zero curve interest rates are 4 bp below the implied interest rates (279 bp vs. 283 bp at 1.4 years and 290 bp vs. 294 bp at 1.9 years, all interest rates are annualized and in logs). To illustrate the resultant effect of interest rates on dividend returns, we use a simple example with index $S = 2,000$, strike price $X = 2,000$, standard deviation $\sigma = 0.2$, dividend yield $\delta = 0.02$, and risk-free rate equal to the implied rate at the required horizon. Based on these parameters, the Black-Scholes call and put prices are 226.63 and 192.70 for the 1.9-year maturity and 193.64 and 171.01 for the 1.4-year maturity. We set the value of collected dividends, $D$, after half a year to 22.00.

From Equation (2) and using the implied interest rates, we find $P^{1.9} = 74.58$ and $P^{1.4} = 55.22$. The corresponding (half-annual) return on the strategy is $(55.22+22.00)/74.58-1 = 0.0354$. Using the zero curve rates instead, we find $P^{1.9} = 73.25$ and $P^{1.4} = 54.04$. The corresponding biased (half-annual) return on the strategy is $(54.04+22.00)/73.25-1 = 0.0381$. Thus, small errors of $-1.26\%$ (at 1.9 years) and $-1.55\%$ (at 1.4 years) in interest rates translate into a sizable error of $+7.63\%$ in the dividend strip return. The resultant average elasticity of the strip return to interest rate errors is large at $-5.49$.

Errors are even larger when we use constant maturity Treasury rates, which are 32 bp and 34 bp below the implied interest rates at 1.4 and 1.9 years. The biased (half-annual) return on the strategy is (46.08+22.00)/63.11 - 1 = 0.0788, which translates into a large error of +122.6%. In light of this magnification of errors in interest rates, we prefer our implied interest rates over any exogenous interest rate.

# Tables and Figures

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

| | Strip ret. | Strip ret. - 2y Treasury ret. | Market ret. | Market ret. - 10y Treasury ret. |
|---|---|---|---|---|
| Mean | 7.63% | 4.34% | 9.02% | 3.91% |
| Std. dev. | 32.83% | 32.85% | 15.78% | 18.57% |
| Sharpe ratio | | 0.13 | | 0.21 |
| AR(1) | -0.34 | -0.34 | 0.01 | 0.06 |
| N | 299 | 299 | 299 | 299 |

Table 1 presents summary statistics for the monthly returns. Returns are continuously compounded (in logarithms of raw returns), annualized, and expressed as a percentage. The period is from January 1996 through December 2020.

Table 2: **Holding Period Returns (Annualized)**

| | 1m | 6m | 12m | 24m | 36m | 48m | 60m |
|---|---|---|---|---|---|---|---|
| Panel A: Strip-return - 2y Treasury return | | | | | | | |
| Mean | 4.34% | 4.67% | 4.26% | 4.54% | 4.56% | 4.45% | 4.51% |
| Std. dev. | 32.85% | 19.28% | 14.64% | 13.28% | 13.32% | 13.20% | 12.92% |
| Sharpe ratio | 0.13 | 0.24 | 0.29 | 0.34 | 0.34 | 0.34 | 0.35 |
| N | 299 | 294 | 288 | 276 | 264 | 252 | 240 |
| Panel B: Market-return - 10y Treasury return | | | | | | | |
| Mean | 3.91% | 3.29% | 2.97% | 2.66% | 2.47% | 1.99% | 1.47% |
| Std. dev. | 18.57% | 19.76% | 19.65% | 20.53% | 19.76% | 18.64% | 17.35% |
| Sharpe ratio | 0.21 | 0.17 | 0.15 | 0.13 | 0.12 | 0.11 | 0.08 |
| Diff. Sharpe ratios ($p$-val.) | [0.73] | [0.68] | [0.38] | [0.02] | [0.00] | [0.00] | [0.00] |
| N | 299 | 294 | 288 | 276 | 264 | 252 | 240 |

Table 2 presents summary statistics for the holding period returns ranging from 1 month through 60 months. Returns are continuously compounded (in logarithms of raw returns), annualized, and expressed as a percentage. In brackets are $p$-values for the HAC test of Ledoit and Wolf (2008) for the difference in Sharpe ratios between Panels A and B. The period is from January 1996 through December 2020.

Table 3: **Predicting the Realized Term Structure**

| | 1m | 6m | 12m | 24m | 36m | 48m | 60m |
|---|---|---|---|---|---|---|---|
| Panel A: | | | | | | | |
| $dp_t^{Strip}$ | 0.26 | 0.59 | 0.71 | 1.02 | 1.11 | 1.01 | 0.79 |
| $t-stat(Overlap.)$ | (4.87) | (4.72) | (3.08) | (2.11) | (2.26) | (1.84) | (1.32) |
| $t-stat(Nonoverlap.)$ | [4.87] | [3.84] | [2.10] | [2.73] | [2.84] | [1.86] | [2.04] |
| $R^2$ | 0.11 | 0.23 | 0.20 | 0.22 | 0.22 | 0.17 | 0.11 |
| Panel B: | | | | | | | |
| $dp_t^{Mkt}$ | -0.01 | -0.12 | -0.30 | -0.64 | -0.85 | -1.04 | -1.10 |
| $t-stat(Overlap.)$ | (-0.39) | (-1.23) | (-1.82) | (-2.59) | (-2.97) | (-3.92) | (-6.42) |
| $t-stat(Nonoverlap.)$ | [-0.39] | [-0.92] | [-1.56] | [-2.42] | [-2.54] | [-4.39] | [-10.88] |
| $R^2$ | 0.00 | 0.03 | 0.11 | 0.26 | 0.38 | 0.53 | 0.65 |
| Panel C: | | | | | | | |
| $dp_t^{Strip} - A^h dp_t^{Mkt}$ | 0.27 | 0.60 | 0.78 | 1.20 | 1.34 | 1.35 | 1.24 |
| $t-stat(Overlap.)$ | (4.86) | (5.65) | (4.54) | (5.16) | (8.12) | (8.61) | (13.26) |
| $t-stat(Nonoverlap.)$ | [4.86] | [4.11] | [3.10] | [4.19] | [4.89] | [5.56] | [10.35] |
| $R^2$ | 0.11 | 0.25 | 0.27 | 0.44 | 0.55 | 0.62 | 0.66 |
| Panel D: | | | | | | | |
| $dp_t^{Strip} - A^h dp_t^{Mkt}$ | 0.40 | 0.84 | 0.96 | 1.37 | 1.43 | 1.47 | 1.35 |
| $t-stat(Overlap.)$ | (5.83) | (8.78) | (5.75) | (6.33) | (9.36) | (10.30) | (19.14) |
| $t-stat(Nonoverlap.)$ | [5.83] | [5.69] | [3.59] | [-4.30] | [5.05] | [11.08] | [34.85] |
| $g_t^{Ind}$ | 0.57 | 1.11 | 0.96 | 1.22 | 0.87 | 1.40 | 1.46 |
| $t-stat(Overlap.)$ | (4.56) | (6.04) | (4.71) | (4.69) | (3.30) | (5.37) | (9.15) |
| $t-stat(Nonoverlap.)$ | (4.56) | (3.78) | (2.51) | (1.90) | (1.26) | (3.69) | (9.32) |
| $R^2$ | 0.16 | 0.33 | 0.32 | 0.48 | 0.57 | 0.67 | 0.72 |

Table 3 presents the results of the predictive regressions for the difference between strip and market returns for holding periods ranging from 1 month through 60 months. In parentheses are t-statistics based on the Newey-West (1987) correction with h lags. In brackets is the average t-statistic across all non-overlapping regressions with different starting months. The period is from January 1996 through December 2020.

Table 4: **Robustness**

| | 1m | 6m | 12m | 24m | 36m | 48m | 60m |
|---|---|---|---|---|---|---|---|
| **Panel A: Strip-return - 2y Treasury return** | | | | | | | |
| Mean | 4.34% | 4.67% | 4.26% | 4.54% | 4.56% | 4.45% | 4.51% |
| Std. dev. | 32.85% | 19.28% | 14.64% | 13.28% | 13.32% | 13.20% | 12.92% |
| Sharpe ratio | 0.13 | 0.24 | 0.29 | 0.34 | 0.34 | 0.34 | 0.35 |
| N | 299 | 294 | 288 | 276 | 264 | 252 | 240 |
| **Panel B: 1.3 year strip-return - 2y Treasury return** | | | | | | | |
| Mean | 4.74% | 5.05% | 4.56% | 4.61% | 4.64% | 4.61% | 4.76% |
| Std. dev. | 33.18% | 18.37% | 14.96% | 13.00% | 13.39% | 13.36% | 12.83% |
| Sharpe ratio | 0.14 | 0.27 | 0.31 | 0.35 | 0.35 | 0.35 | 0.37 |
| N | 299 | 294 | 288 | 276 | 264 | 252 | 240 |
| **Panel C: 0.9 year strip-return - 2y Treasury return** | | | | | | | |
| Mean | 4.91% | 4.99% | 4.53% | 4.69% | 5.02% | 5.00% | 5.12% |
| Std. dev. | 35.81% | 22.94% | 17.67% | 16.28% | 16.29% | 15.81% | 15.68% |
| Sharpe ratio | 0.14 | 0.22 | 0.26 | 0.29 | 0.31 | 0.32 | 0.33 |
| N | 299 | 294 | 288 | 276 | 264 | 252 | 240 |
| **Panel D: 1.9 year zero-curve strip-return - 2y Treasury return** | | | | | | | |
| Mean | 5.96% | 5.57% | 5.19% | 5.57% | 5.71% | 5.64% | 5.89% |
| Std. dev. | 27.98% | 19.88% | 16.80% | 14.28% | 14.22% | 14.11% | 13.52% |
| Sharpe ratio | 0.21 | 0.28 | 0.31 | 0.39 | 0.40 | 0.40 | 0.44 |
| N | 299 | 294 | 288 | 276 | 264 | 252 | 240 |
| **Panel E: 1.9 year strip-return (Moneyness 0.8-1.2) - 2y Treasury return** | | | | | | | |
| Mean | 4.66% | 4.99% | 4.60% | 4.92% | 4.94% | 4.89% | 4.98% |
| Std. dev. | 33.16% | 19.10% | 14.51% | 12.89% | 12.86% | 12.57% | 12.21% |
| Sharpe ratio | 0.14 | 0.26 | 0.32 | 0.38 | 0.38 | 0.39 | 0.41 |
| N | 299 | 294 | 288 | 276 | 264 | 252 | 240 |

Table 4 presents summary statistics for the holding period returns ranging from 1 month through 60 months. Panel A repeats the main strip results from Table 2, Panel A. In Panels B and C, dividend strip returns are based on rolling over investments in strips with maturities of 0.9 year or 1.3 years. In Panel D, dividend prices are estimated using the zero curve interest rate. In Panel D, only options with moneyness in the range 0.8 to 1.2 are used. Returns are continuously compounded (in logarithms of raw returns), annualized, and expressed as a percentage. The period is from January 1996 through December 2020.

Figure 1: **Interest Rates**

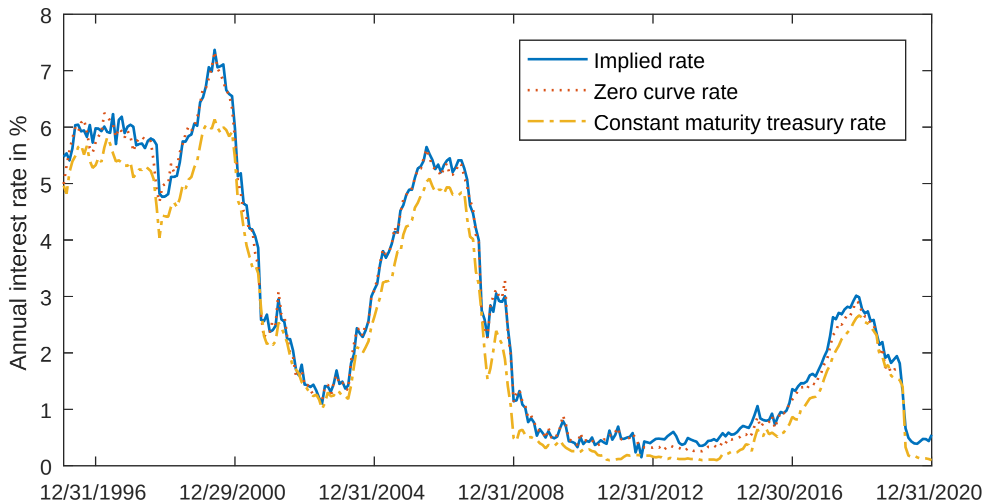


Figure 1 plots the 12-month maturity interest rates. The implied rate is based on put-call parity pairs of S&P 500 index options. The zero curve rate is from OptionMetrics. The constant maturity Treasury rate is from the H.15 filing of the St. Louis Federal Reserve Bank. All interest rates are continuously compounded and expressed as a percentage. The period is from January 1996 through December 2020.

Figure 2: **Cumulative Returns**

Panel A: Cumulative returns

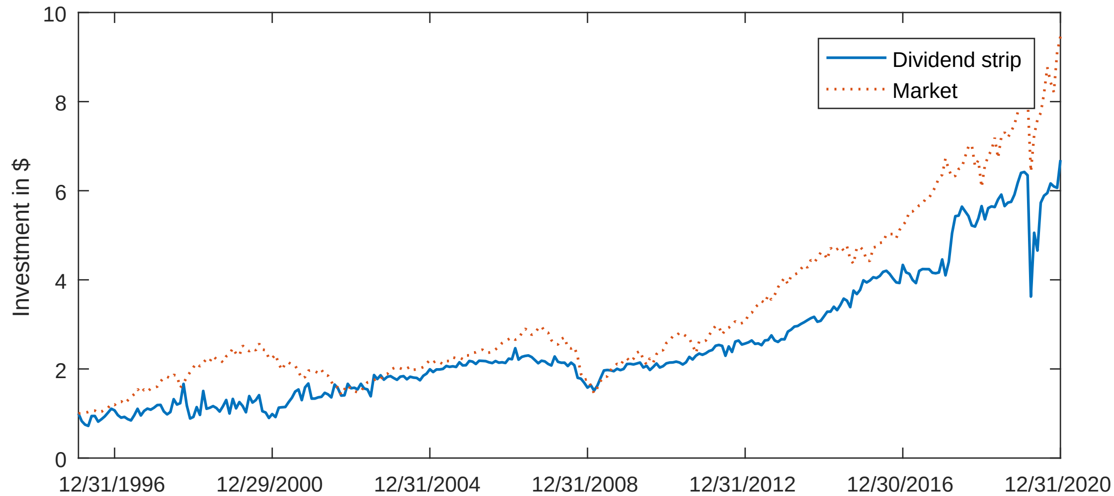


Panel B: Cumulative returns in excess of Treasury returns

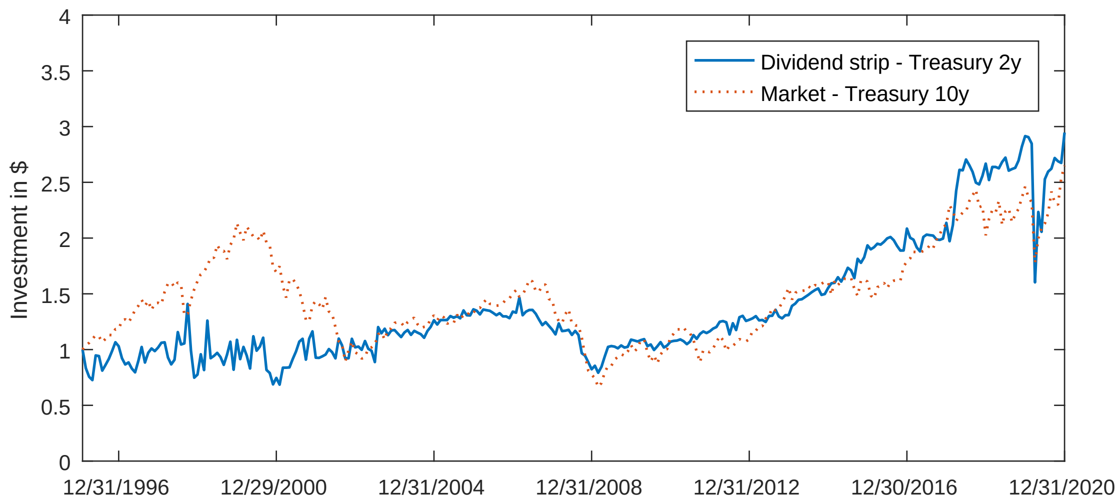


Panel A in Figure 2 plots the cumulative returns for a hypothetical one dollar investment in the dividend strip and the market. Panel B plots the cumulative excess returns of the strip in excess of the two-year Treasury return and the market in excess of the 10-year Treasury return. The period is from January 1996 through December 2020.

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

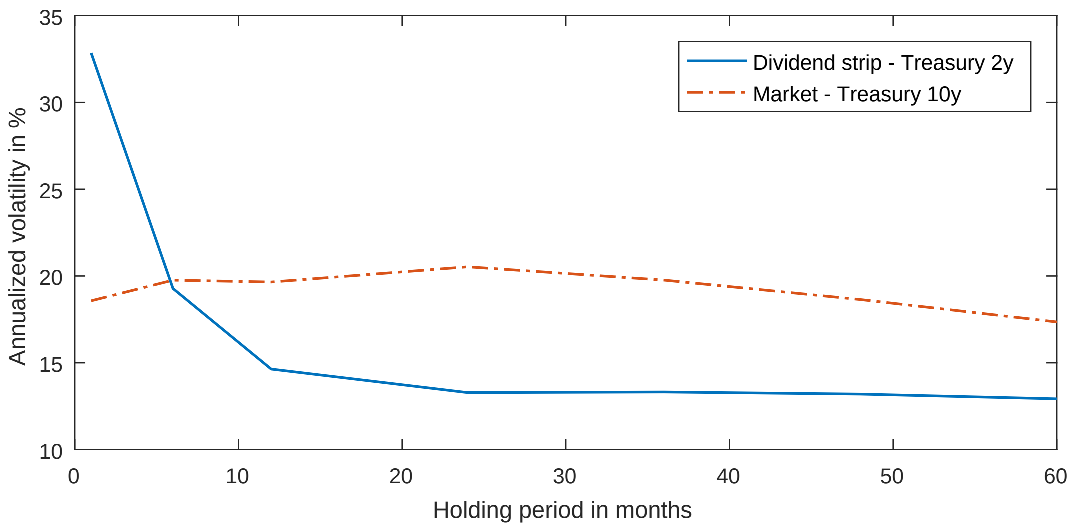


Figure 3 plots the annualized standard deviation for excess strip and market returns for holding periods of 1, 6, 12, 24, 36, 48, and 60 months. The period is from January 1996 through December 2020.

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

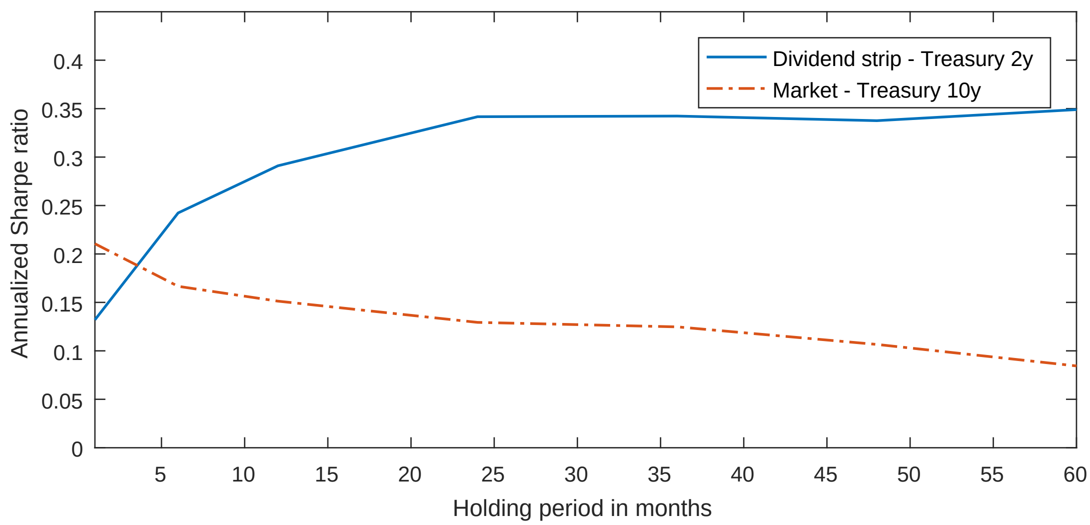


Figure 4 plots the annualized Sharpe ratio for excess strip and market returns for holding periods of 1, 6, 12, 24, 36, 48, and 60 months. The period is from January 1996 through December 2020.

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

Panel A: January 1996 - October 2009

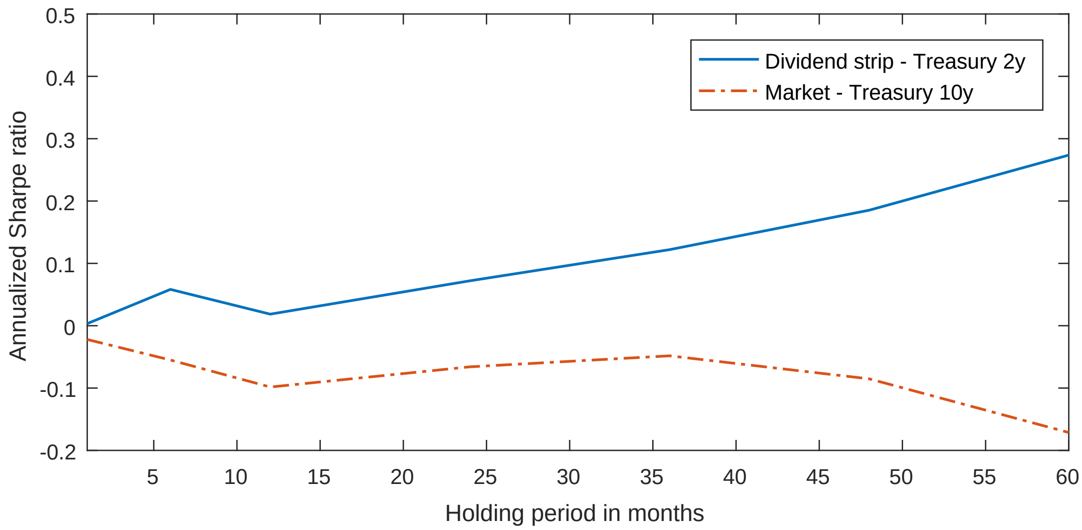


Panel B: December 2004 - December 2020

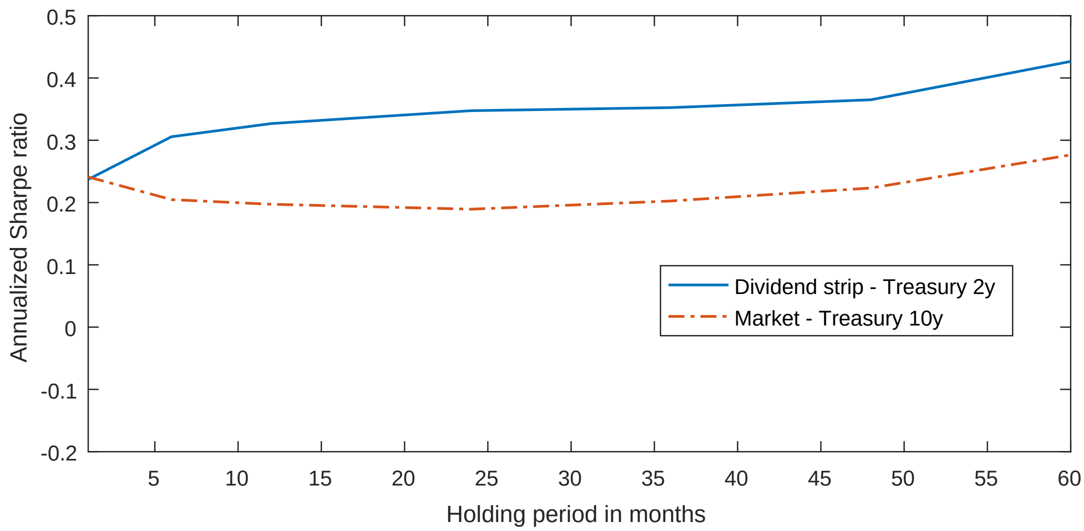


Figure 5 plots the annualized Sharpe ratio for holding periods ranging of 1, 6, 12, 24, 36, 48, and 60 months. The period is from January 1996 through October 2009 (Panel A) and from December 2004 through December 2020 (Panel B).

Figure 6: **Annualized Sharpe Ratios: Business Cycle**

Panel A: Expansions

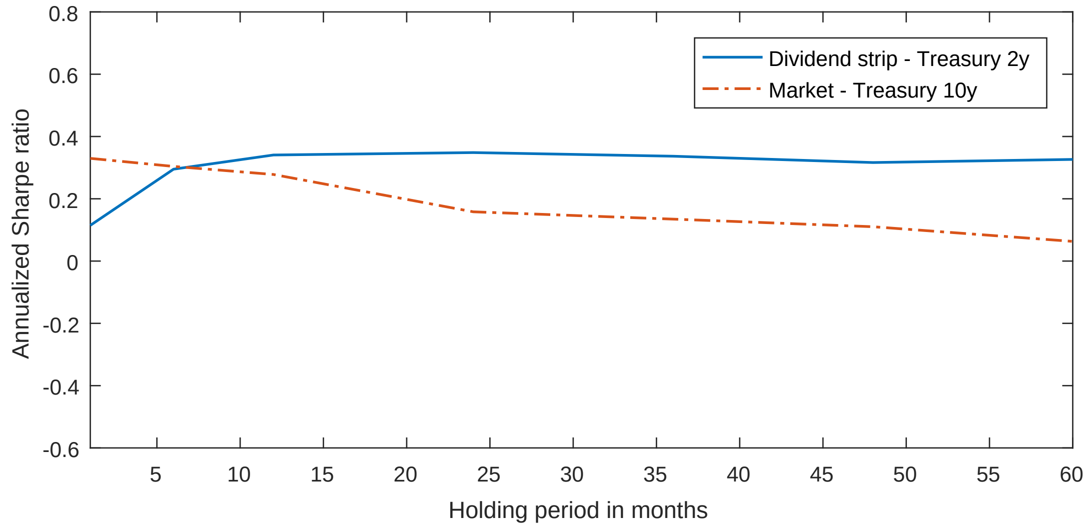


Panel B: Recessions

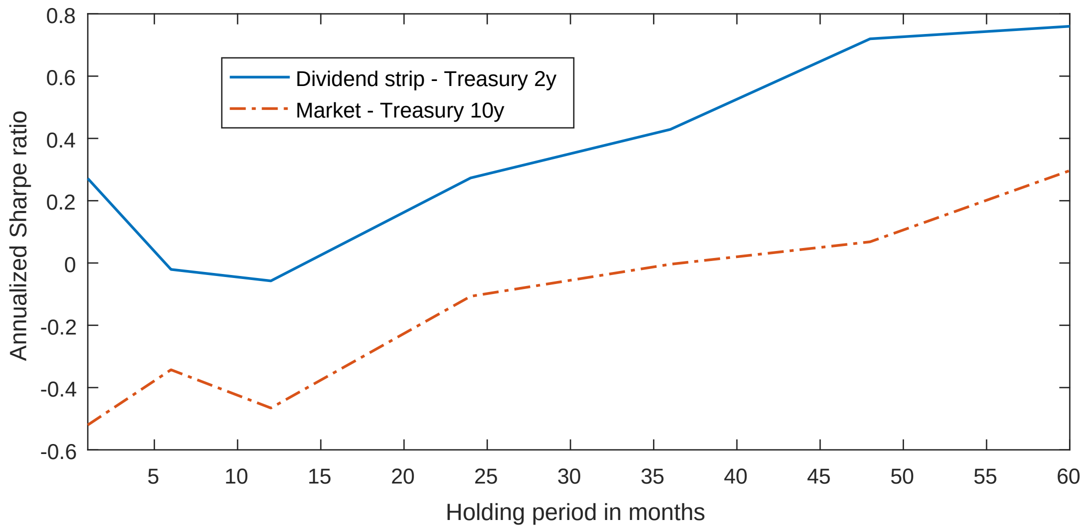


Figure 6 plots the annualized Sharpe ratio for holding periods ranging of 1, 6, 12, 24, 36, 48, and 60 months. The first returns occurs during an NBER expansion (Panel A) and during an NBER recession (Panel B).

# Internet Appendix: A Comparison to Van Binsbergen, Brandt, and Koijen (2012)

In this Internet Appendix, we compare our empirical estimates for dividend strip prices and returns to those reported by Van Binsbergen, Brandt, and Koijen (2012). First, we replicate their estimates. That is, we estimate dividend strip prices from S&P 500 options using the zero curve interest rate. Figure A.1 shows that we are able to almost perfectly replicate their estimates for 12-month dividend strip prices.

Table A.1 reports monthly log returns for dividend strips and the market. We prefer mean logarithmic returns, which are less prone to standard deviation bias. In Panel A, the time period matches the one used by BBK (January 1996 through October 2009). Based on the BBK data, dividend strips (0.85%) outperform the market (0.44%). Our dividend strip returns based on zero curve interest rates closely correlate with the BBK original series (0.98). On average, they are slightly lower than those reported by BBK, but still substantially higher than the market returns, thus confirming BBK's results (0.71% for the strip, 0.44% for the market). When we switch from the zero curve interest rate to the interest rate invariant approach, dividend strips no longer outperform the market (0.41% for the strip, 0.44% for the market).

Next, we extend the sample period through 2020, adding more than a decade of data. Panel B of Table A.1 reports the results. Using the zero curve interest rate during the long sample, we find that dividend strips (0.77%) perform approximately as well as the market (0.75%). Using the interest rate invariant approach during the long sample inverts the relation, and the market (0.75%) outperforms the strip (0.64%).

Table A.1: **Monthly Returns**

| | Strip ret. | Market ret. |
|---|---|---|
| BBK sample (Jan 1996 - Oct 2009) | | |
| Original data | 0.85% | 0.44% |
| Zero curve | 0.71% | 0.44% |
| Interest rate invariant approach | 0.41% | 0.44% |
| Long Sample (Jan 1996 - Dec 2020) | | |
| Zero curve | 0.77% | 0.75% |
| Interest rate invariant approach | 0.64% | 0.75% |

Table A.1 presents the summary statistics for monthly logarithmic returns.

Figure A.1: **Dividend Prices**

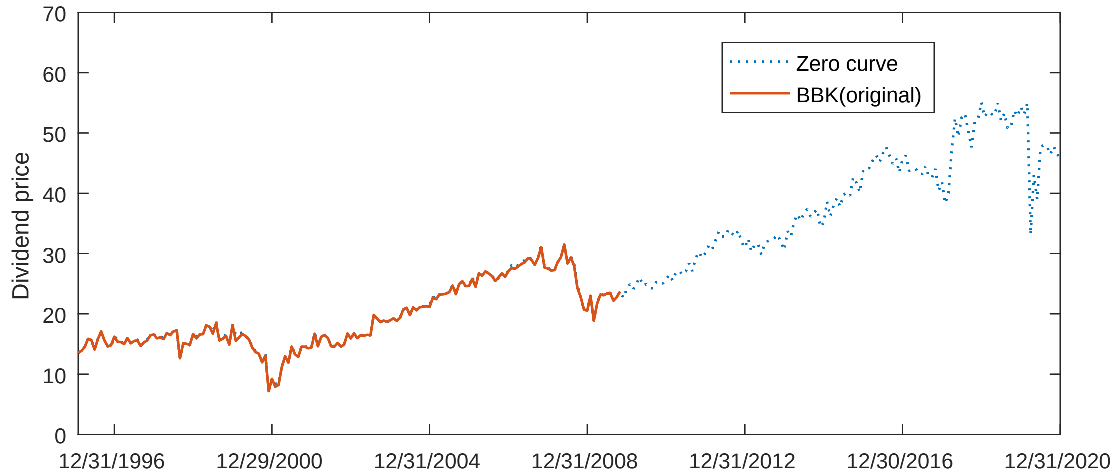


Figure A.1 plots prices for 12-month dividend strips on the S&P 500 index. The period is from January 1996 through December 2020.

Figure A.2: **Annualized Sharpe Ratio Across Different Holding Periods**

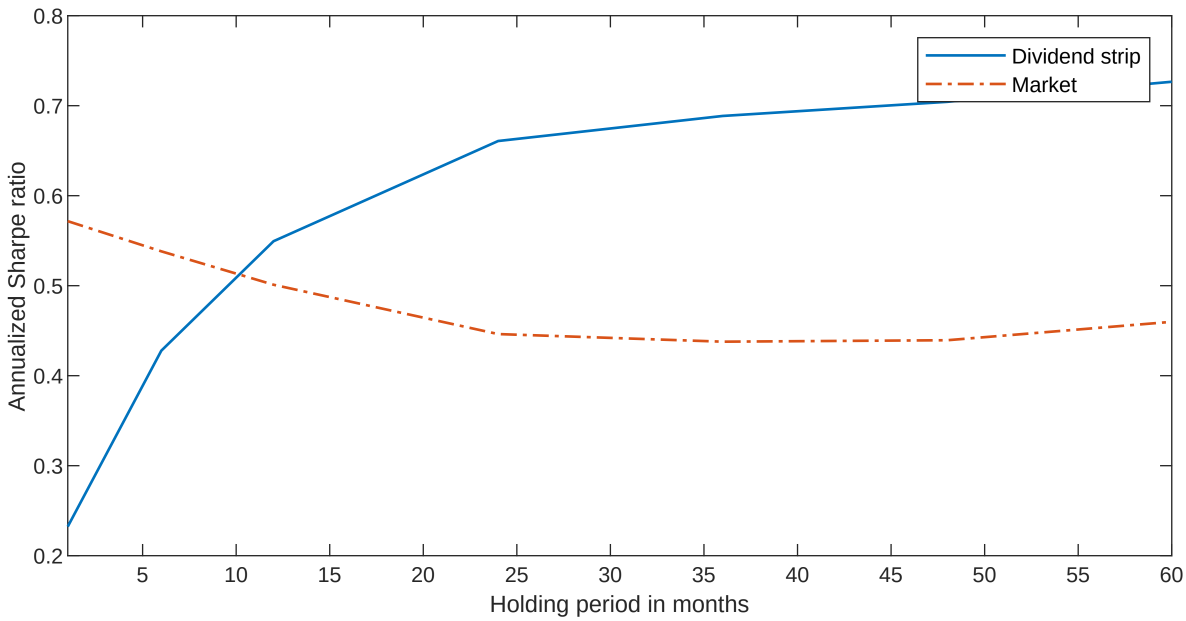


Figure A.2 plots the annualized Sharpe ratio for strip and market returns for holding periods of 1, 6, 12, 24, 36, 48, and 60 months. The period is from January 1996 through December 2020.