Dividend Strips and the Term Structure of Equity Risk Premia: A Case Study of Limits to Arbitrage

Oliver Boguth, Murray Carlson, Adlai Fisher, and Mikhail Simutin *

November 27, 2011

ABSTRACT

Dividend strips, which are claims to dividends paid over future time intervals, shed light on the pricing of risks at different horizons and can help to link asset markets with the real economy. We show, however, that small primary-market pricing frictions become greatly magnified when using standard no-arbitrage assumptions to construct and value synthetic dividend claims. Amplification occurs because of the high implicit leverage of the long-short positions required to replicate dividend strips, and generates remarkably large biases in the moments of synthetic returns. In a calibrated structural model with tiny pricing frictions, we reproduce qualitatively and quantitatively the strongly downward-sloping term structure of equity risk premia, excess volatility, return predictability, and a market beta substantially below one, consistent with empirical evidence. Using more robust return measures we find little support for a statistical or economic difference between the returns to short- versus long-term dividend claims. Our theoretical and empirical analysis shows that no-arbitrage assumptions are not innocuous, particularly for agents wishing to use highly-levered, multi-leg positions to replicate and trade on the relative prices of otherwise illiquid over-the-counter claims.

*Boguth: W. P. Carey School of Business, Arizona State University, PO Box 873906, Tempe, AZ 85287-3906. Carlson and Fisher: Sauder School of Business, University of British Columbia, 2053 Main Mall, Vancouver, BC, V6T 1Z2. Simutin: Rotman School of Management, University of Toronto, 105 St. George Street, Toronto ON, Canada, M5S 3E6. We thank the participants in a brownbag seminar at Arizona State University for helpful early feedback and Lorenzo Garlappi for detailed comments. Support for this project from the Social Sciences and Humanities Research Council of Canada and the UBC Bureau of Asset Management is gratefully acknowledged.

Dividend Strips and the Term Structure of Equity Risk Premia: A Case Study of Limits to Arbitrage

Abstract

Dividend strips, which are claims to dividends paid over future time intervals, shed light on the pricing of risks at different horizons and can help to link asset markets with the real economy. We show, however, that small primary-market pricing frictions become greatly magnified when using standard no-arbitrage assumptions to construct and value synthetic dividend claims. Amplification occurs because of the high implicit leverage of the long-short positions required to replicate dividend strips, and generates remarkably large biases in the moments of synthetic returns. In a calibrated structural model with tiny pricing frictions, we reproduce qualitatively and quantitatively the strongly downward-sloping term structure of equity risk premia, excess volatility, return predictability, and a market beta substantially below one, consistent with empirical evidence. Using more robust return measures we find little support for a statistical or economic difference between the returns to short- versus long-term dividend claims. Our theoretical and empirical analysis shows that no-arbitrage assumptions are not innocuous, particularly for agents wishing to use highly-levered, multi-leg positions to replicate and trade on the relative prices of otherwise illiquid over-the-counter claims.

I. Introduction

Following Mehra and Prescott (1985), numerous authors have attempted to explain the apparently high average returns on equity relative to historical aggregate consumption risk. Two explanations that have received considerable attention are the habit formation model of Campbell and Cochrane (1999) and the long-run risks model of Bansal and Yaron (2004). Both imply under standard calibrations that the long-horizon cash flows of equity are riskier and receive higher returns than short-horizon equity cash flows. In other words, these models generate term structures of equity risk premia that are upward-sloping.[1] Authors such as Lettau and Wachter (2007) point out the importance of the term structure of equity risk premia for evaluating these explanations, and note tensions with other empirical regularities such as the value premium.[2]

Binsbergen, Brandt, and Koijen ("BBK", 2011) propose to measure the term structure of equity risk premia by calculating the returns on dividend strips, which are claims to dividends paid over future time intervals. Using dividend strips one can infer the term structure of equity risk premia; however, dividend strips and dividend swaps are traded in relatively illiquid over-the-counter markets where data are not readily available. The central insight of BBK is that in principle dividend strips can be synthetically replicated by long-short positions in futures and spot markets, relying on futures-spot parity, or by appropriate positions in puts and calls, additionally requiring that put-call parity holds. Using this approach BBK calculate the returns on dividend strips and find higher returns on short-term versus long-term dividend claims. Their evidence contradicts currently prevailing models of the equity premium, which has important implications for theoretical research attempting to understand the link between financial markets and the real economy.[3]

[1]Rare disaster models (Gabaix, 2008, Barro, 2006, Rietz, 1988) imply a flat term structure of equity risk premia. See Binsbergen, Brandt, and Koijen (2011) for a full discussion of the implications of leading asset pricing models for the term structure of equity risk premia, including prior literature and calibrations.

[2]Chen (2011) offers a reconciliation of an upward-sloping equity premium with the value premium by providing alternative measures of cash flow growth. Using these measures, value stocks have longer duration than growth stocks.

[3]Croce, Lettau, and Ludvigson (2011) develop a consumption-based model with a downward-sloping equity premium under imperfect information. Wang (2011) shows that imperfect information can help to match a

We show that using synthetic dividend strips to infer the term structure of equity risk premia can be misleading because of the impact of small pricing frictions. Our claim may seem surprising. BBK emphasize that dividend strips can be replicated using highly liquid futures contracts, and their careful empirical methods and discussion are directly aimed at alleviating concerns about microstructure frictions. The central insight we offer is that, even when pricing frictions have tiny impacts on any leg of a compound trade involving long and short positions, their cumulative effect relative to the value of the net position can be hundreds of times larger. The dramatic amplification of primary-market pricing frictions in a long-short trade occurs because of the effects of implicit leverage.

For example, suppose one buys an index claim and takes a short position in an offsetting one-year futures contract. The net claim is a synthetic dividend strip that entitles one to all dividends paid on the index over the next year. Importantly, due to leverage the net value of the dividend strip is only a small fraction of the gross value of either the long or the short side of the trade. To see this quantitatively, assume an annualized price-dividend ratio of 50, roughly consistent with recent experience.[4] Normalizing the long value to \$100, the offsetting futures position would have a notional amount in the neighborhood of \$98, and the strip value would be about \$2. Now consider that microstructure frictions have tiny impacts of a few basis points (cents) on either the long side or the short side of the trade, or on the synchronicity of the two prices, which is required for parity relations to hold. A few cents of mispricing may be irrelevant compared to the \$100 gross value of the long side, but substantial in comparison with the \$2 net value of the dividend strip. Since measurement errors can impact both the long and short side of the trade, the importance of pricing frictions can easily be one hundred times larger in the dividend strip return than in any leg of the trade in isolation.

downward-sloping equity premium as well as other real and financial market moments in a production economy.

[4]During the 1996-2009 sample period considered by BBK, the U.S. annualized price-dividend ratio averaged approximately 60, ranging from a low of 28 to a high of 90, where the monthly stock market price-dividend ratio is calculated following the method of Shiller (2005) with data downloaded from http://www.econ.yale.edu/ shiller/data.htm.

A symptom of these pricing frictions is the enormous and negative first-order autocorrelation, about -30% at a monthly frequency, of the synthetic dividend returns reported by BBK. Our theoretical analysis shows that two distinct channels can explain these extreme reversals. First, negative autocorrelations are known to be consistent with bid-ask bounce or measurement error (Niederhoffer and Osborne, 1966, Blume and Stambaugh, 1983), and as we show the high implicit leverage of dividend strips dramatically inflates any small measurement errors from the primary markets.

A second and perhaps even more novel result is that asynchronous price adjustment can also contribute substantially to negative autocorrelation in synthetic dividend strip returns. This result is surprising since asynchronous price adjustment is universally associated with positive autocorrelation in portfolio returns (e.g., Lo and MacKinlay (1990)). However, prior literature on asynchronous price adjustment focuses on the effects of asynchronicity in diversified portfolios with positive portfolio weights. For the long-short portfolios we consider the effects are entirely different. In a long-short trade if the long side adjusts to a news or liquidity shock before the short side, the first measured return will miss the hedging effect of the short side, and reversal tends to occur in the next period as the short side catches up. The same reversal logic applies if the short side reacts to information before the long side. This argument thus does not require one side of the trade to always be more informationally efficient than the other. Large negative autocorrelations are generated simply by having the long and the short side react to a shock at not precisely the same time.

Building on Blume and Stambaugh (1983), the negative autocorrelation induced in portfolio returns by microstructure frictions also creates an upward bias in average one-month simple returns. Using a formula provided in Boguth, Carlson, Fisher, and Simutin ("BCFS", 2011), based on the effects of Jensen's inequality, we show that the extreme negative autocorrelation of measured dividend strip returns implies a large average return bias that explains much of the difference between mean short- and long-term asset returns.

Using a theoretical analysis based on log-linearized returns and a calibrated structural model, we show that the principal empirical findings regarding dividend strips can be meaningfully influenced by tiny pricing frictions in primary markets. In a model where the true term structure of equity risk premia is flat, we generate in synthetic dividend strips a strongly downward-sloping term structure, excess volatility, large negative return autocorrelations, return predictability from the price-dividend ratio, and a market model beta substantially less than one. The calibrated model not only qualitatively fits, but also quantitatively matches the key moments.

We further show that alternative measures of returns that are not as sensitive to pricing frictions lead to very different conclusions about the difference between short- and long-term asset returns. Following from BCFS, building on Lo and MacKinlay (1990), while average simple returns can be highly sensitive to microstructure bias at short horizons, average returns over longer horizons are less biased and average logarithmic returns are insensitive to standard microstructure frictions. We find that the apparent higher returns of dividend strips are substantially diminished in annual rather than monthly return intervals, and in average logarithmic returns. We conclude that accounting for the impact of microstructions is important when making inferences regarding the term structure of equity risk premia using synthetic dividend strips.

Our findings provide a stark example of the effects of limits to arbitrage. In constructing dividend strip prices using the standard replicating portfolio approach, one relies on two of the most fundamental no-arbitrage relations in finance, futures-spot and put-call parity. We show that small errors in these relations generate remarkably large effects in the mean, volatility, and predictability of synthetic dividend strip returns. The magnification of seemingly small costs of trading in highly liquid futures and index options markets is an important consideration for agents who may wish to trade on the the relative prices of short- and long-maturity equity risks. By showing the large potential price impacts of small frictions in primary markets, our work adds to the growing literature, following Shleifer and Vishny (1997), on the importance of the limits of arbitrage.

Section II discusses dividend strips and provides a back-of-the-envelope calculation of the bias in average returns due to microstructure effects, following from BCFS. Section III studies microstructure bias in levered portfolios. Section IV gives our calibrated model. Section V provides additional empirical results. Section VI concludes.

II. Dividend Strip Returns

A dividend strip entitles the holder to all dividends paid between dates $t+T_1$ and $t+T_2$, where t denotes the current date. The value of a dividend strip at date t is given by

$$\mathcal{P}_{t,T_1,T_2} \equiv \sum_{\tau=T_1+1}^{T_2} E_t\left(\frac{M_{t+\tau}}{M_t} D_{t+\tau}\right), \tag{1}$$

where M_t and D_t respectively denote the stochastic discount factor and dividend payments at date t. In the special case where $T_1 = 0$, the dividend strip is identical to the "short-term asset" described by BBK, which entitles the holder to dividends paid between the current date t and a future date $t+T_2 = t+T$. Under absence of arbitrage the value of the short-term asset is given by the cost of carry formula for equity futures:

$$\mathcal{P}_{t,T} \equiv \mathcal{P}_{t,0,T} = S_t - e^{-r_{t,T}T} F_{t,T}, \tag{2}$$

where S_t denotes the spot value of the equity claim, $F_{t,T}$ is the forward price, and $r_{t,T}$ is the risk-free rate of interest for a bond maturing at date $t+T$. In the general case where $T_1 \geq 0$ the price of a dividend strip is given by:

$$\mathcal{P}_{t,T_1,T_2} = \mathcal{P}_{t,T_2} - \mathcal{P}_{t,T_1} = e^{-r_{t,T_1}T_1} F_{t,T_1} - e^{-r_{t,T_2}T_2} F_{t,T_2}. \tag{3}$$

In their primary analysis, BBK use put-call parity to rewrite (2) and (3), substituting portfolios of puts, calls, and bonds for futures prices:

$$\mathcal{P}_{t,T} = S_t + p_{t,T} - c_{t,T} - Xe^{-r_{t,T}T}, \tag{4}$$

$$\mathcal{P}_{t,T_1,T_2} = p_{t,T_2} - p_{t,T_1} - c_{t,T_2} + c_{t,T_1} - X(e^{-r_{t,T_2}T_2} - e^{-r_{t,T_1}T_1}), \tag{5}$$

where $p_{t,T}$ and $c_{t,T}$ are respectively puts and calls maturing at $t+T$ with common strike X. BBK use (4) and (5) to calculate prices of dividend strips at the end of each month, using puts and calls on the S&P 500 index, the S&P 500 index value itself, and bond prices.[5] They take great care to attempt to ensure the synchronicity of the various price quotes used to form the synthetic strip prices, and consider two specific return series. The first dividend strip return equals the return on a short-term asset:

$$R_{1,t} = \frac{\mathcal{P}_{t,T} + D_t}{\mathcal{P}_{t-1,T+1}}, \tag{6}$$

where the maturity T varies, depending on the derivatives traded, between approximately 1.3 and 1.9 years for each month of implementation.[6] The second dividend strip return is

$$R_{2,t} = \frac{\mathcal{P}_{t,T_1,T_2}}{\mathcal{P}_{t-1,T_1+1,T_2+1}}, \tag{7}$$

where T_2 is set identically equal to the maturity T used in the first strategy, and $T_1 \approx T_2 - 12$ is chosen to match the available contract with maturity approximately one year earlier. The second strategy does not require replicating the index or collecting dividends, and involves trades only in highly liquid futures contracts.

A. Properties of the Return Series

BBK report that over a February 1996 to October 2009 sample period, the return series $R_{1,t}$ and $R_{2,t}$ are highly correlated. They provide the following key facts:

- The one-month average returns of both short-term dividend strips are substantially larger than the one-month average returns on the S&P 500 (annualized 11.6% and 11.2% versus 5.6%). The return differences persist after controlling for standard risk factors.

- The volatilities of the short-term dividend strips are substantially higher than the S&P 500 index (standard deviations of monthly returns are 7.8% and 9.7%

[5] As a robustness check, BBK also calculate strip prices using futures and compare with the prices based on puts and calls. We comment on the validity of this robustness check later in our paper.

[6] Specifically, in January of any given year, the maturity T is chosen according to the available contract expiring in the fall of the following year. This contract will have a maturity of about $T = 1.85$ at purchase. This contract is held for six months, during which period T decreases by 1/12 each month. On July 1, this contract is sold and a new contract with maturity of approximately $T = 1.85$ years is purchased.

versus 4.7%). The volatilities of the dividend strips are substantially larger than the volatilities of subsequent dividend realizations, which is diagnosed as excess volatility.

- The dividend strip returns have estimated betas of about 0.5 in market model regressions.
- The returns series $R_{1,t}$ is highly predictable by lagged values of the price-dividend ratio of a 1.5 year short-term dividend strip.

An important indicator of measurement error in the BBK results is extreme negative autocorrelation, about -0.30, in each of the two return series. Clearly, the return series $R_{1,t}$ and $R_{2,t}$ are not tradeable at the assumed prices. If they were, an average return of 30% per month per dollar invested would be available by following a return reversion investment strategy. Hence, some degree of measurement error must be present in the dividend strip return series.

B. Measurement Error and Bias in Average Returns

The effect of measurement error on inferences about expected returns and performance is a fundamental topic in the finance literature. Early literature such as Blume and Stambaugh (1983) focuses on the effects of microstructure frictions on daily rebalanced portfolios, including equal-weighted strategies. This research concludes that measurement error induces negative autocorrelation in returns and inflates measured daily returns. Roll (1983) investigates the effects of daily rebalancing over multiple periods. Recently, Asparouhova, Bessembinder, and Kalcheva (2011) show that even lower-frequency monthly rebalancing can result in strong biases in returns, especially for portfolios of illiquid stocks whose average returns can be overstated by more than 0.4% monthly.

The effects of microstructure bias on the measured returns of buy-and-hold portfolios, such as an investment in the value-weighted S&P index, has historically focused on autocorrelations and volatilities (Niederhoffer and Osborne, 1966, Scholes and Williams, 1977). For example, it is well-known that the S&P index has a significantly positive first order autocorrelation at a one-month horizon. This positive autocorrelation in the index return is generally understood to be largely driven by asynchronous adjustment in the index component prices, and not an indication of true trading profits from a short-term persistence strategy, as shown for example in Boudoukh, Richardson, and Whitelaw (1994), Ahn, Boudoukh, Richardson, and Whitelaw (2002).

Despite the significant impacts of microstructure frictions on the second moments of returns, the classical literature has largely perceived microstructure frictions to be unproblematic for calculations of average returns in the absence of periodic rebalancing. This is likely due to the seminal work of Scholes and Williams (1977) and Lo and MacKinlay (1990), who show, among other contributions, that standard microstructure frictions do not alter mean returns. However, both of these prior studies consider only logarithmic returns.

In recent work, BCFS show that microstructure frictions impact not only the autocorrelations and volatilities of buy-and-hold portfolios, as shown in prior literature, but also their average simple returns and alphas. For example, under positive autocorrelations high observed short-horizon returns tend to follow one another, raising the benefit of compounding and increasing the long-horizon return. To compensate, the average short-horizon return must be downward biased. BCFS show that average monthly returns of long-only buy-and-hold stock portfolios can be impacted by as much as 0.25% per month relative to longer-horizon averages. Most empirical studies that compare returns of different investment strategies, including BBK, focus on average simple monthly returns.

C. Approximating the Impact on Dividend Strip Average Returns

We use a formula provided in BCFS to estimate the return a buy-and-hold investor can expect to obtain from following a strategy of trading in the short-term asset. Let $r_{it} = \ln(R_{it}) \sim \mathcal{N}(\mu_i, \sigma_i^2)$ denote a series of monthly log returns. We further assume that log returns aggregated over n months have a normal distribution with variance

σ_{in}^2, where n is the relevant horizon. These assumptions are exactly satisfied if r_{it} is a stationary ARMA(p, q) process with Gaussian innovations, and approximately hold in more general cases. The buy-and-hold mean is $\bar{R}_{in}^{BH} \equiv \mathbb{E}\left[R_{i,t+1} \cdots R_{i,t+n}\right]$ and we consider the corresponding rescaled monthly measure defined by $\bar{R}_{in}^{RS} \equiv \mathbb{E}\left[R_{it}\right]^n$. BCFS show that the relative difference between the two return measures satisfies

$$\frac{\bar{R}_{in}^{RS}}{\bar{R}_{in}^{BH}} = e^{n\sigma_i^2(1-VR_{in})/2}, \tag{8}$$

where $VR_{in} \equiv \sigma_{in}^2 / \left(n\sigma_i^2\right)$ is the variance ratio.

The bias in rescaled monthly returns relative to the buy-and-hold return is thus determined by the monthly variance σ_i^2 and return autocorrelations as summarized by the variance ratio VR_{in}. When the variance ratio is one, for example if returns are iid, rescaled monthly returns are an unbiased estimate of the investment return of the buy-and-hold investor. For assets with positive return autocorrelations, variance ratios typically exceed one and equation (8) suggests that their buy-and-hold return exceeds the compounded short horizon return. A negative autocorrelation implies high short- relative to long-horizon returns. All else constant, the bias is more severe when σ_i^2 is high. Although the formula (8) is based on the assumption of lognormal distributions, BCFS show that it is empirically extremely accurate for a wide range of portfolios.

Table 1 shows our calculation of the ratio (8) for dividend strip returns and the S&P500, using moments reported by BBK. Panel A lists the primary moments used in the formula. Both of the dividend strip return series $R_{1,t}$ and $R_{2,t}$ exhibit large volatilities and very strong negative autocorrelations. These moments suggest that average one-month returns rescaled to an annual horizon will be substantially upward biased relative to the actual twelve-month average return. In comparison, the S&P500 return series shows moderate positive autocorrelation, lower average returns, and lower volatility.

Panel B presents the average one-month returns for each series rescaled to an annual horizon, the estimated twelve-month buy-and-hold return average, and an estimate of the variance ratio for each investment.[7] For the S&P index, the average monthly return rescaled to an annual horizon is 6.93%. The variance ratio of 1.16 exceeds one due to positive autocorrelation in the return series, and the estimated annual buy-and-hold return of 7.16% slightly exceeds the rescaled monthly return. By contrast, the average monthly returns of the dividend strip series when rescaled to an annual horizon appear to suggest an annual return exceeding 14%. However, with variance ratios far below one, 0.51 and 0.33 for R_1 and R_2 respectively, the estimated annual buy-and-hold returns of 12.8% and 10.1% are substantially lower than the average monthly returns would suggest. The bias in the short-horizon monthly returns is inferred to be about 2% annually for R_1 and 4.2% annually for R_2.

These first calculations show that the extreme negative autocorrelations of dividend strip returns and apparent bias in short-horizon monthly returns relative to longer-horizon annual returns are related symptoms. The root causes of these symptoms can be traced to small pricing frictions, as we now show.

III. Microstructure Bias in Levered Portfolios

It is not obvious that microstructure frictions should have a large impact on the average monthly returns of synthetic dividend strips. After all, as noted by BBK the synthetic dividend strips can be created by investments in highly liquid derivatives of the S&P index. The impact of microstructure frictions on any one of these contracts should be small. The key idea of our paper, which we now develop more formally, is that the high leverage implicit in the synthetic dividend strip strategy magnifies negligible microstructure frictions in any leg of the strategy.

We assume that the observed index level S_t^o is equal to the sum of the true unob-

[7] Since no variance ratios are reported in BBK, we use the 1-lag autocorrelation to approximate the variance ratio using

$$VR(q) = 1 + 2\sum_{k=1}^{q-1} (1 - k/q)\, \rho_k, \tag{9}$$

as in Campbell, Lo, and MacKinlay (1997, Eq. 2.4.19). For the higher order autocorrelations, we assume $\rho_k = 0$ for $k \geq 2$. We can alternatively impose the restrictions implied by an AR process, specifying $\rho_k = \rho_1^k$ for $k \geq 2$. The estimated horizon effects are affected little by this assumption.

servable index level S_t and two error terms:

$$S_t^o = S_t + \rho(S_{t-1} - S_t) + S_t(e^{\eta_t} - 1). \tag{10}$$

The term $\rho(S_{t-1} - S_t)$, $0 \leq \rho \leq 1$, accounts for lead-lag effects among the index constituents. The term $S_t(e^{\eta_t} - 1)$ captures iid measurement error. The futures price is for now taken to be unaffected by microstructure frictions (i.e., $F_{t,T}^o = F_{t,T}$), and we weaken this assumption in the next section. If one uses the futures-spot parity relation (2) to impute the short-term asset value, the observed price is

$$\mathcal{P}_{t,T}^o \equiv S_t^o - e^{-r_{t,T}T} F_{t,T}^o \tag{11}$$

$$= \mathcal{P}_{t,T} + (S_t^o - S_t). \tag{12}$$

This equation shows that the effect of microstructure friction on the short-term asset is identical in absolute terms to the effect on the index. The effect is proportionally larger in the short-term asset, however, since $\mathcal{P}_{t,T}$ is considerably smaller than S_t.

Using a procedure similar to Campbell and Shiller (1988), we now provide log-linear approximations for the observed prices of the index and short-term asset. Let lower case letters denote the logarithm of their uppercase counterpart, and define the capital gain return $R_t^x = S_t/S_{t-1}$. This allows us to equivalently express the observed index level in equation (10) as a proportion of its true value,

$$S_t^o = S_t \left[\rho \left(\frac{1}{R_t^x} - 1 \right) + e^{\eta_t} \right]. \tag{13}$$

Let $\ell_{t,T} \equiv \ln(S_t/\mathcal{P}_{t,T})$ be the logarithm of the implicit leverage of the dividend strip. Also define $\bar{\ell}_T = \mathbb{E}(\ell_{t,T})$ and $\bar{L}_T = exp(\bar{\ell}_T)$. In the special case where $\ell_{t,T}$ is a constant, $\bar{L}_T$ is the implicit leverage of the dividend strip with maturity T. We then show

Proposition 1 *Logarithmic index levels and short-term asset prices are:*

$$s_t^o \approx s_t - \rho r_t^x + \eta_t, \tag{14}$$

$$p_{t,T}^o \approx s_t - \ell_{t,T} + \bar{L}_T(-\rho r_t^x + \eta_t). \tag{15}$$

Microstructure frictions have larger effects on the log price of the short-term asset, and the amplification is summarized by $\bar{L}_T$, reflecting leverage. For dividend strips with a one-year maturity, $\bar{L}_T^{-1} \approx 0.02$ approximates the dividend yield of the index, and $\bar{L}_T \approx 50$ approximates the price-dividend ratio of the index. Pricing frictions that are small relative to the observed index level (14) can therefore be dramatically magnified in the measured dividend strip price (15).

Equations (14) and (15) also show that when the dividend strip leverage $\ell_{t,T}$ is stationary in t holding T constant, the observed log prices s_t^o and $p_{t,T}^o$ are cointegrated, sharing the common stochastic trend s_t. This feature is a natural consequence of the short-term nature of the pricing frictions we assume. Over long time periods, we therefore expect the prices of the index and the dividend strip to grow at approximately the same rate. Since long-run growth rates are determined by geometric mean returns, the geometric means of the capital gains returns of both series should be similar.[8] The analysis therefore suggests that if over long periods of time total-return indices of the index and the dividend strip are to differ, it must be through the dividend yield – as opposed to the capital gain – component of returns. That is, if the term structure of equity risk premia is not flat, it must be reflected over the long run in different averages of the log dividend yields of the stock market index and the dividend strip.

The closed-form expressions for prices in Proposition 1 allow direct analysis of the impact on returns. Let $\Delta_t = \ln(D_t/S_t)$, $\Delta = E(\Delta_t)$, $\delta_t = \ln(D_t/\mathcal{P}_{t,T})$, and $\delta = E(\delta_t)$. We then show

Proposition 2 *Observed logarithmic returns on the index and the short-term asset are respectively*

$$r_t^o \equiv \frac{S_t^o + D_t}{S_{t-1}^o} \quad \approx \quad r_t + \rho(r_{t-1}^x - r_t^x) + \eta_t - \eta_{t-1}, \tag{16}$$

$$r_{1t}^o \equiv \frac{\mathcal{P}_{t,T}^o + D_t}{\mathcal{P}_{t-1,T+1}^o} \quad \approx \quad r_{1t} + \bar{L}_T[\rho(r_{t-1}^x - r_t^x) + \eta_t - \eta_{t-1}] \tag{17}$$

[8]Average log capital gains returns should also be similar since the average log return is a simple transform of the geometric mean return.

where

$$r_t \equiv \frac{S_t + D_t}{S_{t-1}} \approx \ln(1+e^{\Delta}) + \frac{e^{\Delta}}{1+e^{\Delta}}(\Delta_t - \Delta) + s_t - s_{t-1}, \tag{18}$$

$$r_{1t} \equiv \frac{\mathcal{P}_{t,T} + D_t}{\mathcal{P}_{t-1,T+1}} \approx \ln(1+e^{\delta}) + \frac{e^{\delta}}{1+e^{\delta}}(\delta_t - \delta) + p_{t,T} - p_{t-1,T+1}. \tag{19}$$

Equations (16) and (17) show that the observed log index return has the same mean as the true log index return, and the observed dividend strip log return has the same mean as the true dividend strip log return. Consistent with prior literature, neither form of microstructure friction introduces bias into the observed log returns. However, any small microstructure frictions do become greatly magnified in the logarithm of the short-term asset return. The size of the amplification effect is driven by the implicit leverage $\bar{L}_T$ of the synthetic dividend strip.

Despite the fact that mean log returns are not biased by microstructure frictions, average simple returns can be greatly affected. Following from the standard Jensen's inequality approximation:

$$\mathbb{E}(R_{it}) \approx e^{E(r_{it})+0.5\mathrm{Var}(r_{it})}, \tag{20}$$

any microstructure-induced bias in volatity should impact simple returns. In recent work, BCFS show that (20) is empirically extremely accurate for a broad set of style portfolios, in the sense that for a range of measurement horizons, the scaling of simple returns is largely explained by the behavior of measured volatilities. We therefore anticipate bias in average simple returns to be driven by the magnitude of any volatility bias created by microstructure frictions.

We now show that the impact of microstructure frictions on the measured variance of the dividend strip can be large. For tractability, consider the case where dividend yields on the index and dividend strip are constants, and the true index return r_t is iid. The variance of the short-term asset return is then

$$\mathrm{Var}(r^o_{1t}) = \mathrm{Var}(r_t)\left[1 - 2\rho\bar{L}_T + 2\rho^2\bar{L}_T^2 + 2\bar{L}_T^2\frac{\mathrm{Var}(\eta_t)}{\mathrm{Var}(r_t)}\right]. \tag{21}$$

Lead-lag effects inflate variance if $\rho > \bar{L}_T^{-1}$, and measurement error unambiguously increases volatility. Both effects are multiplied by factors on the order of $2\bar{L}_T^2$. If we again consider $\bar{L}_T = 50$, implicit leverage magnifies microstructure-induced variance by 5000 times! Even small primary-market pricing frictions cause apparent excess volatility, which biases the average simple return following (20).

Consistent with standard implications of excess volatility (Shiller, 1981), the observed short-term asset return is predictable with first-order autocovariance

$$\mathrm{Cov}(r^o_{1t}, r^o_{1t-1}) = \mathrm{Var}(r_t)\left[\rho\bar{L}_T - \rho^2\bar{L}^2_T - \bar{L}^2_T\frac{\mathrm{Var}(\eta_t)}{\mathrm{Var}(r_t)}\right]. \tag{22}$$

Lead-lag effects can cause positive or negative autocovariance depending on whether $\bar{L}^{-1}_t - \rho$ is positive or negative, while measurement error unambiguously biases the measured autocovariance downwards. Leverage again inflates the importance of both frictions by factors on the order of $\bar{L}^2_T$.

The microstructure frictions also impact measured beta. The covariance of the short-term asset return with the index return is

$$\mathrm{Cov}(r^o_t, r^o_{1t}) = \mathrm{Var}(r_t)\left[1 - \rho(1+\bar{L}_T) + 2\rho^2\bar{L}_T + 2\bar{L}_T\frac{\mathrm{Var}(\eta_t)}{\mathrm{Var}(r_t)}\right]. \tag{23}$$

Lead-lag effects reduce the measured covariance when $1 + \bar{L}^{-1} > 2\rho > 0$. Monthly S&P 500 returns are positively autocorrelated with $\rho \approx 0.1 < (1 + \bar{L}^{-1}_T)/2$, and should therefore contribute to downward bias in covariance. The last term in (23) shows that measurement error increases covariance. The magnitude depends on the variance of the measurement error relative to the variance of index return, again magnified by the implicit leverage $\bar{L}_T$ of the short-term asset. The observed beta

$$\beta^o_{1t} \equiv \mathrm{Cov}(r^o_t, r^o_{1t})/\mathrm{Var}(r^o_t) \tag{24}$$

is determined by dividing (23) by the measured index variance

$$\mathrm{Var}(r^o_t) = \mathrm{Var}(r_t)\left[1 - 2\rho + 2\rho^2 + 2\frac{\mathrm{Var}(\eta_t)}{\mathrm{Var}(r_t)}\right]. \tag{25}$$

Since leverage does not affect the index variance, the absolute bias in (25) is expected to be tiny relative to the absolute bias in the covariance (23). The effect of microstructure frictions on the observed dividend strip beta (24) should therefore be dominated by bias in the covariance (23).

Table 2 uses the equations derived in this section to illustrate the impact of implicit leverage and microstructure frictions on various moments of the short-term asset return. The selection of parameters is consistent with the ratio $\mathcal{P}/S$ varying from 1/10 to 1/90. Variation in this parameter can either reflect different assumptions for the annual dividend yield on the S&P 500 index or consideration of a variety of dividend strip strategies with $T - t$ ranging from a few months to over a year. Lead-lag effects in the index are captured by ρ, roughly equal to the first-order autocorrelation of index returns. Finally, the parameter $\sigma(\eta)$ reflects the magnitude of measurement errors.

The table shows effects of microstructure frictions on the monthly mean return, standard deviation, autocorrelations, and market beta. Measurement error biases measured volatility upwards, in extreme cases more than tripling the unobserved true level of 4%. Average simple returns are also overstated, in some cases more than doubling their true value. Even modest measurement error creates substantial negative autocorrelation. Observed market betas are only modestly inflated. Asynchronous price adjustment also has large impacts. At leverage levels of $\bar{L}_T$ equal to 30 or higher, even small values of ρ can significantly increase measured standard deviations and average simple returns. Negative autocorrelations are easily produced by high leverage. At low leverage ($\bar{L}_T = 10$), the bias in standard deviations and simple return averages is negative, and return autocorrelations are positive. Significant downward bias in beta occurs for all tabulated levels of $\bar{L}_T$ and ρ. The table shows that by combining the effects of measurement error and asynchronous price adjustment, all of the key facts associated with dividend strips can be produced.

IV. A Structural Model, Calibration, and Implications

We now show that in a simple calibrated model, tiny microstructure effects capture qualitatively and quantitatively the primary empirical features of dividend strips. Let dividends X_t be given by

$$dX_t = gX_t + \sigma X_t dW_t, \tag{26}$$

where dW_t is the increment of a Wiener process, g is the mean growth rate, and σ^2 the variance. The fair price of the equity claim is given by the Gordon growth formula:

$$S_t = \frac{X_t}{\mu - g}, \tag{27}$$

where μ is the constant return on equity.

A. Delayed Price Adjustment in the Index

We first consider the case where only the observed level of the index at time t, denoted S_t^o, depends on lagged prices:

$$S_t^o = (1 - \rho_1 - \rho_2)S_t + \rho_1 S_{t-1} + \rho_2 S_{t-2}. \tag{28}$$

True and measured returns are given by:

$$R_{Mt} = S_t/S_{t-1} - 1, \tag{29}$$

$$R_{Mt}^o = S_t^o/S_{t-1}^o - 1. \tag{30}$$

The true value of a short-term asset with a claim on the first T years of dividends starting from date t is given by:

$$\mathcal{P}_{t,T} = S_t(1 - e^{-(\mu-g)T}). \tag{31}$$

We initially assume that the observed futures price is based on the fair value of the index:

$$F_{t,T}^o = F_{t,T} = (S_t - \mathcal{P}_{t,T})\, e^{rT}. \tag{32}$$

The observed price of the short-term asset, based on futures prices and the observed index value, is then

$$\mathcal{P}_{t,T}^o = \mathcal{P}_{t,T} + \rho_1(S_{t-1} - S_t) + \rho_2(S_{t-2} - S_t). \tag{33}$$

In the special case where $\rho_2 = 0$, the measured return on the short-term asset is

$$\frac{\mathcal{P}_{t+1,T-1}^o}{\mathcal{P}_{t,T}^o} = \frac{\mathcal{P}_{t+1,T-1} - \rho_1(S_{t+1} - S_t)}{\mathcal{P}_{t,T} - \rho_1(S_t - S_{t-1})}. \tag{34}$$

As in the lognormal approximation of Section 3, the presence of asynchronous trading drives negative autocorrelation in the returns of the short-term asset. This is opposite to the effect of asynchronous trading in long-only portfolios, which causes positive autocorrelations (e.g., Lo and MacKinlay, 1990). Additionally, the bias caused by slow price adjustment negatively correlates with the true price change, and positively correlates with the lagged true price change. These effects will cause a downward bias in the estimated market beta.

Table 3, Panel A, shows the magnitudes of these effects. To calibrate the model, we set $r_f = 0.0029$ per month and $\mu = 0.0056$ equal to the sample averages of the monthly risk-free rate and market return reported by BBK in their sample. We consider two values for the growth rate g of dividends. First, in column (i) we set $g = 0.042$, which is high from a historical perspective but necessary to approximately match the price-dividend ratio of about 60 over this sample period. Alternative calibrations that use a lower dividend growth rate as well as a lower risk-free rate or market risk premium to approximately match the observed price-dividend ratio produce results similar to those reported in this calibration. Second, in column (ii) we choose a more conservative value $g = 0.025$ for the growth rate of dividends, which produces a much smaller value for the aggregate price-dividend ratio of about 27. In both calibrations, we choose $\sigma = 0.047$ to approximately match monthly market volatility.

In column (i) with a high price-dividend ratio of 60 the average return, autocorrelation, and volatility of the short-term asset are greatly affected by small amounts of asynchronous price updating in the market index. We set $\rho_1 = 0.03$ and $\rho_2 = 0.01$, which gives a first-order autocorrelation of the market index of 0.0313, less than half of what is observed empirically during this sample. This small amount of asynchronous price reaction produces a measured return of the short-term asset of 1.10% per month, approximately equal to the average return of the strategy R_1 reported by BBK. The first-order autocorrelation of the short-term asset produced by the model is -0.2582, also very close to the empirical value of -0.2682. The model also produces large excess volatility of the short-term asset. The simulated standard deviation of short-term asset returns is 0.1160, substantially larger than the volatility of the market return.

Table 3 also reports the market beta from a regression of the short-term asset return on a constant, the market return, and the lagged short-term asset return. In the absence of asynchronous prices, the market beta coefficient in this regression would be equal to one. As suggested by Section 3, the presence of delayed price reaction in the index can substantially bias downward the estimated beta in the market model regression. In this calibration the downward bias in beta is much stronger than observed in the data, with the simulated market model beta of -1.44 substantially less than the empirical value of 0.448. In untabulated results, a regression of the observed short-term asset return on lagged values of the short-term asset price-dividend ratio also produces a significantly negative coefficient (-0.58 in the simulation versus -0.17 empirically). Overall, these results show that a simple model of asynchronous price adjustment in the market index can quantitatively match the mean return and autocorrelation of market and short-term asset returns, while also qualitatively capturing excess volatility in the short-term asset, downward bias in the market model beta, and negative predictability of short-term asset returns from the short-term asset price-dividend ratio.

In column (ii) of Panel A we carry out a similar analysis with a more modest growth rate of dividends that produces an aggregate price-dividend ratio of about 27, much smaller than observed during the BBK sample. To compensate for the smaller price-dividend ratio, we set $\rho_1 = 0.075$ and $\rho_2 = 0.015$, which are larger than the first calibration. These values produce a first-order autocorrelation of market returns of 0.0822, almost exactly matching the empirically observed statistic of 0.0898. The simulated moments of the dividend strip returns are very similar in this calibration to those in column (i). The average returns of the short-term asset are slightly higher (0.0122) and the autocorrelations somewhat more negative (-0.3277). Both calibrations show the ability to match important quantitative and qualitative aspects of short-term asset returns, simply by introducing small amounts of asynchronicity in the market index.

Replicating the short-term asset return R_1 considered in this subsection requires trading in the spot market to replicate the market index. The costs of such a trading strategy may not be small. For this reason, BBK suggest that a more appropriate strategy to empirically evaluate is the dividend steepener return given in (7), which we now consider.

B. Measurement Error in Prices and the Dividend Steepener

In principle, the dividend steepener return in (7) can be obtained by trading only in futures markets, with a long position in maturity T_1 and a short position in maturity $T_2 > T_1$. BBK choose $T_2 \approx 18$ months and $T_1 \approx T_2 - 12$ months. In practice, most of the analysis in BBK is based not on trading in futures markets, but on approximate replication of the steepener strategy by trading in puts and calls of equivalent maturities and relying on the accuracy of the put-call parity relations (4) and (5).

In this subsection, we consider the impact of small amounts of measurement error in the futures prices imputed from market data on the calculated average returns of the short-term asset. We assume no asynchronicity of prices in this subsection, but rather permit that observed futures prices relate to fair value by

$$F^o_{t,T} = F_{t,T} e^{\eta_{t,T}}. \tag{35}$$

We assume for convenience that the $\eta_{t,T}$ are independently drawn from a normal distribution with mean zero and standard deviation σ_η. Hence, the observed prices are unbiased estimates of the true prices, but may contain small errors. Consistent with the idea of limits to arbitrage, our setup permits that standard no-arbitrage relations hold very closely, but not exactly.

In column (i) of Panel B in Table 3, we again set the growth rate of dividends $g = 0.0042$ to approximately match the average aggregate price-dividend ratio of the market of about 60 during this sample period. We then set $\sigma_\eta = 0.0009$, which reflects a fairly small measurement error. For comparison, the average bid-ask spread in liquid futures contracts during this period was about 3 basis points, which gives an idea of the potential accuracy of price measurements in actual futures market prices. Moreover, the prices used by BBK in most of their analysis are not actual futures prices, but are imputed from put-call parity using quotes matched as closely as possible in time from S&P 500 index options markets. While BBK make every effort to be as accurate as possible in their approach, the reality of limits to arbitrage and imperfect observability of wholly synchronous prices in real time makes it plausible to assume small measurement errors in prices.[9]

The consequences of these small measurement errors are shown in Table 3. In column (i) of Panel B the average observed return of the dividend steepener is 1.17% per month, more than double the true dividend steepener return, which equals the true market return of 0.056% per month. Hence, small, zero-mean errors in observed prices produce a large bias in the calculation of average returns for the dividend steepener. Other symptoms of the small measurement errors in prices are the large negative autocorrelation of returns of the measured dividend steepener returns, equal to -0.4165. This negative autocorrelation is entirely the consequence of small measurement errors, as the autocorrelation of the true steepener return series is zero by construction. The steepener return series also displays substantial excess volatility.

Similar effects can be seen in column (ii) of Panel B, where we more conservatively set $g = 0.0025$ to generate a much lower level of the price-dividend ratio, about 27. In this case, we increase the measurement error standard deviation to $\sigma_\eta = 0.0015$. The effect on average returns and variances is smaller than column (i), but still economically meaningful, inflating mean returns of the dividend steepener by about 50% relative to their true value (0.0086 versus 0.0056).

Hence, either small amounts of price asynchronicity (Panel A of Table 3) or measurement error (Panel B) can generate high returns, strong negative autocorrelation, and excess volatility of short-term dividend strips as observed in the data. The primary

[9]For evidence of the imperfect correspondence between observed futures prices and those imputed from put-call parity, see Figure 5 of BBK.

difference in the consequences of these two microstructure frictions can be seen in the last row of Table 3, which shows that the market model beta of the dividend steepener return is approximately one. Hence, while price asynchronicity causes a downward bias in the market model beta coefficients of dividend strip returns (Panel A), iid measurement error does not cause substantial bias in the estimated market model regression coefficient (Panel B). The empirical values of the market model regression betas are 0.448 for R_1 and 0.4863 for R_2, in between the model values in Panels A and B. This suggests that a combination of price asynchronicity and measurement error will be able to approximately match all of the main features of dividend strip returns shown in Table 3, as we now show.

C. Full Model

Panel C of Table 3 shows two ways of combining price asynchronicity with measurement error. In the first case, we set

$$F_{t,T}^{o} = [(1-\rho_{1,T})F_{t,T} + \rho_{1,T}F_{t-1,T}]e^{\eta_{t,T}}.$$

This permits a small amount of slow price adjustment in the futures price, in addition to measurement error. In column (i) we show results for a calibration with a tiny amount of slow price adjustment in the 6-month contract that forms the long-side of the steepener contract ($\rho_{1,T_1} = 0.009$) with no price delays in the 18-month contract that forms the short side of the steepener. Combined with measurement error, this specification matches the mean returns, autocorrelations, excess volatility, and market model regression beta of the empirical data.

Alternatively, it seems unlikely that any one futures contract is always more informationally efficient than another liquid contract, in the sense of always incorporating current pricing information sooner and more fully. To accommodate that information about prices may at different times enter more quickly into one contract or the other, we allow that an iid draw of a random variable determines which contract receives a news shock at each date t. Specifically, let observed futures prices be given by

$$F^o_{t,T} = \begin{cases} F_{t,T} e^{\eta_{t,T}} & \text{if } \mathbf{1}_{t,T,lagger} = 0 \\ \left[(1-\rho_{1,lagger})F_{t,T} + \rho_{1,T} F_{t-1,T}\right] e^{\eta_{t,T}} & \text{if } \mathbf{1}_{t,T,lagger} = 1, \end{cases}$$

where $\mathbf{1}_{t,T,lagger}$ takes the value zero when the contract with maturity T fully incorporates available information at time t, and takes the value one when the contract includes lagged price information. We assume that at each date t, one of the contracts T_1 and T_2 used to form the steepener fully incorporates all available information, and the other is a lagger. The realization of which contract is the lagger is an iid draw in each period, and in our calibration we set $P\left([\mathbf{1}_{t,T_1,lagger} = 1] = 2/3\right)$, so that the 18 month contract on average has more current information. This specification, shown in column (ii) of Panel C also allows us to match the mean returns, autocorrelations, excess volatility, and market model regression beta of the empirical data on observed dividend strip returns. Hence, small amounts of asynchronicity and measurement error combine to produce the principal features of dividend strip returns.

D. Implicit Leverage and Average Returns

The previous sections have shown that mean-zero microstructure frictions, either asynchronous price adjustment or measurement error, cause observed short-term dividend strips to have upward-biased mean returns, apparent excess volatility, and negative correlation in returns. We now show how these effects change as we vary the horizon of a short-term dividend strip, which affects the implicit leverage of the long-short position.

Figure 2 shows the implicit leverage, average returns, volatility, and autocorrelation of the dividend strip given by the short-term asset $\mathcal{P}^o_{t,T}$, as we vary the maturity T from 6 to 60 months. All other aspects of the calibration in this example are identical to Table 3, Panel A, column (i). As the maturity of the dividend strip shortens, the implicit leverage in the long-short position increases dramatically. Assuming an annualized price-dividend ratio of approximately 60, the average over the BBK sample period, a six-month dividend strip has implicit leverage exceeding 100. Small microstructure frictions in primary markets become greatly magnified in the synthetic return series, as reflected in the high returns and excess volatility shown in the figure at short maturities. For longer maturities the implicit leverage falls and the importance of the microstructure effects in the synthetic return series lessens.

In the third panel, the autocorrelation function shows a distinctive pattern of first being negative, reaching a minimum at a short maturity less than one year, and then increasing to eventually become positive. The changing autocorrelation reflects a changing balance between two effects. On the one hand, the value of the short-term asset is a fraction of the value of the aggregate market, and the observed returns of the aggregate market have positive persistence. When the dividend strip has a long maturity this effect dominates. On the other hand, the appearance in the numerator and denominator of (34) of the current and lagged true price change causes negative autocorrelation in returns. This effect is more pronounced for more levered dividend strips with shorter maturities T, consistent with our results in Section 3.

Figure 3 shows similar plots in the case where returns are impacted by iid measurement errors rather than asynchronous price changes. Again, the implicit leverage of dividend strips magnifies the importance of microstructure effects, and shorter maturities show higher apparent average returns and excess volatility. The autocorrelations in the third panel are always negative consistent with our results in Section 3, and become closer to zero for longer maturities.

In both figures, synthetically replicating a short-term dividend strip requires high leverage in nearly offsetting long and short positions. The greater the leverage, the more the amplification of small pricing errors in the fundamental securities used to create the replicating portfolio. The true term structures of risk premia and volatilities underlying both figures are flat. Hence, the amplification of small microstructure frictions entirely generates the apparent downward-sloping term structures of risk premia and volatilities in both figures.

E. Cointegration and Price versus Return Correlations

Most of the empirical analysis in BBK is based on using the put-call parity relations (4) and (5) to impute futures prices from option markets. The only analysis in BBK using actual futures prices is a paragraph on robustness (their Section 5.2), which shows that the correlations of short-term asset *prices* calculated from options versus futures markets are 94% and 91% for 6-month and 1-year maturities respectively. BBK conclude that because futures markets are highly liquid and the short-term asset prices obtained from futures and options markets are similar, microstructure frictions are an unlikely explanation for their findings.

It is important to distinguish between high correlation of prices and high correlation of returns. In all of the models we have developed in this paper, the correlation between fair value and observed short-term asset prices is high. Nonetheless, we find large differences in average returns. Figure 4 shows this using a simulation of 180 months of prices and returns from the model of Table 3, Panel C, column (ii). The correlation of prices in this example is 93.67%. Hence the implied measurement error in true versus actual returns in the model is of a similar magnitude to the implied measurement error in the difference between actual futures prices and futures prices inferred from options markets in the results reported by BBK (see also their Figure 5). However, despite the high correlation of the two price series in the model the correlation of the measured and true monthly returns are only 19.82%.

The large difference in the autocorrelations of prices and returns is possible because the true and actual return series in the model are cointegrated. Attempts by arbitrageurs to take advantage of profit opportunities imply that the difference between true and measured prices in financial markets must be small and temporary. However, returns may be different and are predictable in the short-run by differences between true and measured prices. A similar cointegrating relationship holds for futures prices measured directly or via put-call parity from the options markets. In either case, the similarity of prices does not imply a high degree of closeness in returns, either realized or on average. In particular, despite the high price correlation between true and measured prices of the short-term asset shown in Figure 4, the average monthly measured return is twice as high as the average monthly true return (1.12% versus 0.56%). The higher average of the measured return is offset by its strong negative autocorrelation, which substantially dampens the positive effects of compounding over longer horizons, following our discussion in Section 2. In general, high correlation of two price series, or apparent cointegration between two price series, is not informative about the closeness of their average returns. In particular when one series is subject to strong microstructure bias, generating significant return autocorrelations, the short-horizon average must be biased to compensate. These biases are greatly mitigated in returns measured over longer horizons, as we now show.

V. New Estimates of the Dividend Term Premium

In this Section, we reexamine the performance of dividend strip strategies using return measures that are more robust to microstructure frictions. We find no evidence that short-term assets outperform the index by a statistically significant margin.

We extract returns of the two investment strategies (R_1 and R_2) from BBK (2010).[10] Table 4 shows that moments from the extracted data precisely match those reported in BBK.[11] Monthly returns of both investment strategies are extremely volatile and show strong negative autocorrelation. Consistent with the arguments in Section 2, average monthly returns are therefore biased upward. To confirm this, Table 5 compares averages of returns compounded at different horizons. Whereas the average monthly return of the dividend steepener strategy (R_2) reaches 1.13% (14.44% per year), the average annual return scaled to a monthly frequency amounts to just 0.72% (8.99% per year). By contrast, the average returns of the S&P 500 index are relatively similar across

[10]Figures 5 and 6 of the BBK (2010) working paper are plotted using vector graphics, which allows us to accurately extract monthly returns using Adobe Photoshop.

[11]We use the value-weighted S&P 500 index from CRSP. Market excess return, HML and SMB factors, and the risk-free rate are from Ken French's data library, http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

horizons, equalling 0.58% and 0.53% respectively for monthly returns and annual returns rescaled to a monthly horizon. Consequently, using longer-horizon annual returns where the impact of microstructure frictions are proportionately smaller, the difference in performance between the dividend steepener strategy and the index reaches only a statistically insignificant 0.19% monthly.

Most of the difference in average returns of the short-term asset and the index is attributable to the first half of the sample. For example, Table 5 shows that the difference in average returns R_1 of the short-term asset and the S&P 500 index amounts to 0.90% monthly in the first half but reaches just 0.23% in the second half. Autocorrelation of the short-term asset returns is more negative and volatility is higher during the 1996-2002 period. Moreover, liquidity of the assets used to compute returns R_1 and R_2 was likely lower earlier in the sample. It is thus not surprising that the horizon effects are particularly striking in the first half of the sample. For example, when the annual compounding horizon is used, the difference in average returns of the short-term asset R_1 and the index amounts to 0.49% monthly in the first half and 0.30% in the second half of the sample. More striking, the corresponding difference in returns of the dividend steepener strategy R_2 and the index is *negative*, measuring -0.05% per month during the first half of the sample, remarkably lower than the monthly return mean difference of 0.72% that BBK focus on.

Figure 4 graphically confirms that the cumulative performance of the short-term asset and the index are very similar. An investment in the dividend steepener strategy in 1996 produced a lower cumulative return than an investment in the S&P 500 index through 2002. A one dollar investment in either strategy grew to the same amount ($2.10) by the end of June 2004. Moreover, during the last four years of the sample (2006-2009), cumulative returns of the two investments were largely similar. Thus the difference in average returns of the dividend steepener (1.13%) and the index (0.58%) are not representative of what a long-term investor would realize. To illustrate this point further, Figure 4 also plots cumulative returns of two hypothetical strategies that every month earn 1.13% and 0.58%. Comparing the differences, the one-month average returns give a remarkably misleading view of the performance of the strategy. These biases are entirely consistent with the effects of small pricing errors in primary markets, which become amplified when using no-arbitrage relations to impute fair values of dividend strips.

VI. Conclusion

We show that small mean-zero pricing frictions can create enormous biases in the measured short-horizon average returns, autocorrelations, volatilities, and predictability of synthetic dividend strips. Dramatic amplification of small microstructure effects occurs because of the leverage implicit in the synthetic dividend strips, which are created from long-short positions in derivative markets and typically have a net value only a small fraction of the long and short positions used in their construction. Using calibrated models where fundamental returns are iid and the term structure of equity risk premia is flat, but permitting small mean-zero frictions, we are able to match all of the principal empirical features of measured short-horizon synthetic dividend strip returns.

Empirically, longer-horizon returns are less impacted by short-run microstructure effects than monthly returns, and present a quite different picture of the returns to a dividend strip investment strategy. The cumulative returns to the dividend steepener strategy are actually below the cumulative returns of an investment in the market index over nearly the entire first half of the 1996-2009 sample, while the average monthly returns appear to show a sixty basis point per month outperformance over the same period. Over the entire sample period, we find little evidence of a statistical or economic difference in the longer-horizon returns to dividend strip strategies.

Our results show the importance of limits to arbitrage even for seemingly reliable relationships such as futures-spot and put-call parity. Small errors in these relationships can become greatly magnified when one constructs long-short positions to attempt to create return profiles that are not directly tradeable in financial markets. These biases particularly impact short-horizon average returns, variances, autocorrelations, and beta estimates. We expect there may be other interesting applications of these ideas, for example in evaluating the performance of pairs trading strategies, or in the evaluation of hedge fund strategies that require high leverage.

Appendix

Proof of Proposition 1

Equation (14) follows from a Taylor series expansion of the two variable equation (13) around $r_t^x = \ln(R_t^x) = 0$ and $\eta_t = 0$. Equation (15) follows from a Taylor series expansion of the expression

$$\begin{aligned} \mathcal{P}^o_{t,T} &= \mathcal{P}_{t,T} + (S_t^o - S_t) \quad (36) \\ &= S_t(e^{-\ell_{t,T}} + \rho e^{-r_t^x} - (1+\rho) + e^{\eta_t}) \quad (37) \end{aligned}$$

around the three points $\ell_{t,T} = \bar{\ell}_T$, $r_t^x = 0$, and $\eta_t = 0$. ∎

Proof of Proposition 2

Equations (18) and (19) are the standard Campbell and Shiller (1988) log-linearizations of the index and dividend strip returns. Equations (16) and (17) then follow from first-differencing the relevant expressions for log ex-dividend index and dividend strip prices in equations (14) and (15). Note that to simplify the expression, the expansion for $\mathcal{P}^o_{t-1,T+1}$ is around $\bar{\ell}_T$, which represents the average dividend-price ratio for the short-term asset $\mathcal{P}^o_{t,T}$. ∎

References

Ahn, Dong-Hyun, Jacob Boudoukh, Matthew P. Richardson, and Robert F. Whitelaw, 2002, Partial adjustment or stale prices? implications from stock index and futures return autocorrelations, *Review of Financial Studies* 15, 655–689.

Asparouhova, Elena, Hendrik Bessembinder, and Ivalina Kalcheva, 2011, Noisy prices and inference regarding returns, Working Paper, University of Utah, UT.

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

Barro, Robert J., 2006, Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics* 121, 823–866 067CX Times Cited:46 Cited References Count:36.

Binsbergen, Jules H. van, Michael W. Brandt, and Ralph S. J. Koijen, 2011, On the timing and pricing of dividends, *American Economic Review* , Forthcoming.

Blume, Marshall E., and Robert F. Stambaugh, 1983, Biases in computed returns - an application to the size effect, *Journal of Financial Economics* 12, 387–404.

Boguth, Oliver, Murray Carlson, Adlai J. Fisher, and Mikhail Simutin, 2011, On horizon effects and microstructure bias in average returns and alphas, Working Paper.

Boudoukh, Jacob, M. P. Richardson, and Robert F. Whitelaw, 1994, A tale of three schools: Insights on autocorrelations of short-horizon stock returns, *Review of Financial Studies* 7, 539–573.

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, 205–251.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press: Princeton, New Jersey).

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

Chen, Jason, 2011, What does the value premium tell us about the term structure of equity returns?, Working Paper.

Croce, Massimiliano, Martin Lettau, and Sydney Ludvigson, 2011, Investor information, long-run risk, and the duration of risky cash-flows, Working Paper, UC Berkeley, CA.

Gabaix, Xavier, 2008, Variable rare disasters: A tractable thoeory of ten puzzles in macro-finance, *American Economic Review* 98, 64–67 308EE Times Cited:0 Cited References Count:13.

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, 55–92.

Lo, Andrew W., and A. Craig MacKinlay, 1990, An econometric analysis of nonsynchronous trading, *Journal of Econometrics* 45, 181–211.

Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.

Niederhoffer, Victor, and M. F. M. Osborne, 1966, Market making and reversal on the stock exchange, *Journal of the American Statistical Association* 61, 897–916.

Rietz, Thomas A., 1988, The equity risk premium - a solution, *Journal of Monetary Economics* 22, 117–131.

Roll, Richard, 1983, On computing mean returns and the small firm premium, *Journal of Financial Economics* 12, 371–386.

Scholes, Myron S., and Joseph Williams, 1977, Estimating betas from nonsynchronous data, *Journal of Financial Economics* 5, 309–327.

Shiller, Robert J., 1981, Do stock prices move too much to be justified by subsequent changes in dividends, *American Economic Review* 71, 421–436.

——— , 2005, *Irrational Exuberance* (Princeton University Press: Princeton, New Jersey) 2nd edition edn.

Shleifer, Andrei, and Robert W. Vishny, 1997, The limits of arbitrage, *Journal of Finance* 52, 35–55.

Wang, Shubo, 2011, Incomplete information and asset returns in real business cycle economies, Working Paper, University of British Columbia, Canada.

Table 1. Back of the Envelope Calculation

	A. Estimates from BBK One-month Returns			B. 12-month Returns		
Portfolio	$\bar{R}$	σ	ρ_1	$\bar{R}_{12}^{RS}$	VR_{12}	$\bar{R}_{12}^{BH}$
$S\&P$	0.0056	0.0469	0.0898	0.0693	1.1646	0.0716
R_1	0.0116	0.0780	-0.2682	0.1484	0.5083	0.1280
R_2	0.0112	0.0965	-0.3668	0.1430	0.3275	0.1008
Difference						
$R_1 - S\&P$	0.060	-0.3580	0.0791	0.0791		0.0564
$R_2 - S\&P$	0.056	-0.4570	0.0737	0.0737		0.0292

Notes: This table reports in Panel A estimated moments for the $S\&P$ index and the two short-term assets from Binsbergen, Brandt, and Koijen (2011). Panel B compares the rescaled simple monthly return $\bar{R}_{in}^{RS} \equiv \mathbb{E}[R_{it}]^n$ with an estimated buy-and-hold return $\bar{R}_{in}^{BH} \equiv \mathbb{E}[R_{i,t+1} \cdots R_{i,t+n}]$, $n = 12$ obtained from the approximation

$$\frac{\bar{R}_{in}^{RS}}{\bar{R}_{in}^{BH}} = e^{n\sigma_i^2(1-VR_{in})/2},$$

provided in Boguth, Carlson, Fisher, and Simutin (2011). The key parameter in the approximation, the variance ratio $VR_{in} \equiv \sigma_{in}^2 / \left(n\sigma_i^2\right)$ is computed under the assumption that all autocorrelations of order greater than one are zero.

Table 2. Microstructure Bias in Leveraged Portfolios

$\bar{L}_T$	$\rho = 0$			$\rho = 0.05$			$\rho = 0.1$		
	$\sigma(\eta) = 0$	0.05%	0.1%	$\sigma(\eta) = 0$	0.05%	0.1%	$\sigma(\eta) = 0$	0.05%	0.1%
A. Return Standard Deviation (%)									
10	4.00	4.06	4.24	2.83	2.92	3.16	4.00	4.06	4.24
30	4.00	4.53	5.83	6.32	6.67	7.62	14.42	14.58	15.03
50	4.00	5.34	8.12	11.66	12.19	13.64	25.61	25.86	26.57
70	4.00	6.36	10.68	17.20	17.90	19.85	36.88	37.21	38.18
90	4.00	7.52	13.34	22.80	23.67	26.12	48.17	48.58	49.82
B. Expected Simple Return (%)									
10	0.30	0.30	0.31	0.26	0.26	0.27	0.30	0.30	0.31
30	0.30	0.32	0.39	0.42	0.44	0.51	1.27	1.29	1.36
50	0.30	0.36	0.55	0.90	0.97	1.16	3.56	3.63	3.82
70	0.30	0.42	0.79	1.71	1.84	2.21	7.27	7.40	7.80
90	0.30	0.50	1.12	2.86	3.07	3.70	12.55	12.78	13.46
C. Return Autocorrelation									
10	0.00	-0.02	-0.06	0.50	0.44	0.30	0.00	-0.02	-0.06
30	0.00	-0.11	-0.26	-0.30	-0.32	-0.36	-0.46	-0.46	-0.46
50	0.00	-0.22	-0.38	-0.44	-0.45	-0.46	-0.49	-0.49	-0.49
70	0.00	-0.30	-0.43	-0.47	-0.48	-0.48	-0.49	-0.49	-0.49
90	0.00	-0.36	-0.46	-0.48	-0.49	-0.49	-0.50	-0.50	-0.50
D. Market Beta									
10	1.00	1.00	1.01	0.55	0.56	0.57	0.12	0.13	0.14
30	1.00	1.01	1.04	-0.44	-0.43	-0.40	-1.83	-1.82	-1.78
50	1.00	1.02	1.06	-1.44	-1.42	-1.37	-3.78	-3.76	-3.70
70	1.00	1.02	1.09	-2.43	-2.41	-2.33	-5.73	-5.70	-5.62
90	1.00	1.03	1.11	-3.43	-3.39	-3.30	-7.68	-7.65	-7.53

Notes: This table reports moments of the ex-dividend return of the short-term asset for varying microstructure parametrization ($\rho \in \{0, 0.05, 0.1\}$, $\sigma(\eta) \in \{0, 0.0005, 0.001\}$) and for different price-dividend ratios ($\bar{L}_T \in \{10, 30, 50, 70, 90\}$) following the approximations from Proposition 2. Panel A shows return standard deviations, Panel B expected simple ex-dividend returns, Panel C return autocorrelations, and Panel D observed market betas. The log ex-dividend return has a mean of 0.22% and a standard deviation of 4%, so that simple return in the absence of microstructure frictions average 0.3%.

Table 3. COMPARISON OF MODEL VERSUS EMPIRICAL MOMENTS

	A. Asynchronous Prices		B. Measurement Error		C. Full Model		Empirical Value	
Parameter	(i)	(ii)	(i)	(ii)	(i)	(ii)	R_1	R_2
T_1	0	0	6	6	6	6	0	6
T_2	12	12	18	18	18	18	12	18
g_d	0.0042	0.0025	0.0042	0.0025	0.0042	0.0042		
ρ_{1,T_1}	0.03	0.075	0	0	0.009	0		
ρ_{2,T_1}	0.01	0.015	0	0	0	0		
σ_η	0	0	0.0009	0.0015	0.0009	0.0003		
$\rho_{1,lagger}$	-	-	-	-	-	0.0025		
$P(lagger = T_1)$	-	-	-	-	-	2/3		
Moment								
$\bar{R}_M$	0.0056	0.0056	0.0056	0.0056	0.0056	0.0056	-	-
$\rho_1(R_M)$	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	-	-
$\bar{R}_\mathcal{P}$	0.0056	0.0056	0.0056	0.0056	0.0056	0.0056	-	-
$\rho_1(R_\mathcal{P})$	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	-	-
$\bar{R}^o_M$	0.0054	0.0054	-	-	-	-	0.0056	-
$\bar{R}^o_M - \bar{r}_f$	0.0025	0.0025	-	-	-	-	0.0027	-
$\rho_1(R^o_M)$	0.0313	0.0822	-	-	-	-	0.0898	-
$std(R^o_M)$	0.0454	0.0431	-	-	-	-	0.0469	-
$PD(R^o_M)$	60	27	60	27	60	60	60	-
$\bar{R}^o_\mathcal{P}$	0.0110	0.0122	0.0117	0.0086	0.0110	0.0112	0.0116	0.0112
$\rho_1(R^o_\mathcal{P})$	-0.2582	-0.3277	-0.4165	-0.3703	-0.4088	-0.4080	-0.2682	-0.3668
$std(R^o_\mathcal{P})$	0.1161	0.1249	0.1201	0.0947	0.1156	0.1157	0.0780	0.0965
$\beta^o_\mathcal{P}$	-1.4401	-1.4635	1.0010	1.0030	0.4731	0.4767	0.4480	0.4863

Notes: This table shows moments of observed market and dividend strip returns under calibrations of the models described in Section 3. All models set the risk-free rate to $r_f = 0.029$ per month, the market return drift to $\mu = 0.056$, and the variance of dividends to $\sigma = 0.047$ per month. In each column, the model is simulated for 200,000 months to obtain the moments reported in the table.

Table 4. COMPARISON OF EXTRACTED AND BBK MOMENTS

	R_1		R_2		S&P 500	
	BBK	Extracted Data	BBK	Extracted Data	BBK	CRSP VW Index
A. Full Sample						
Mean	0.0116	0.0115	0.0112	0.0113	0.0056	0.0058
Median	0.0079	0.0081	0.0148	0.0148	0.0106	0.0109
Std. Dev.	0.0780	0.0781	0.0965	0.0963	0.0469	0.0468
Sharpe ratio	0.1124	0.1106	0.0872	0.0876	0.0586	0.0635
AR(1)	-0.2682	-0.2719	-0.3668	-0.3825	0.0898	0.0887
Market alpha	0.0073	0.0071	0.0069	0.0069		
Market beta	0.4721	0.4729	0.4847	0.4835		
Market model R^2	0.0877	0.0878	0.0604	0.0603		
FF alpha	0.0065	0.0064	0.0053	0.0053		
FF market beta	0.4880	0.4891	0.5712	0.5700		
FF value beta	0.1393	0.1396	0.3744	0.3769		
FF size beta	0.0751	0.0739	-0.0279	-0.0253		
FF model R^2	0.0915	0.0916	0.0811	0.0813		
B. First Half						
Mean	0.0159	0.0158	0.0139	0.0140	0.0065	0.0068
Median	0.0117	0.0113	0.0231	0.0243	0.0093	0.0096
Std. Dev.	0.0986	0.0987	0.1212	0.1208	0.0514	0.0514
Sharpe ratio	0.1242	0.1230	0.0843	0.0856	0.0456	0.0619
C. Second Half						
Mean	0.0072	0.0071	0.0086	0.0085	0.0046	0.0047
Median	0.0058	0.0058	0.0086	0.0087	0.0118	0.0120
Std. Dev.	0.0494	0.0495	0.0630	0.0630	0.0422	0.0418
Sharpe ratio	0.1060	0.1027	0.1044	0.1033	0.0615	0.0658

Notes: This table reports moments of the two short-term assets (R_1 and R_2) and the S&P 500 index. BBK moments are taken from Binsbergen, Brandt, and Koijen (2011). To obtain Extracted Data moments, R_1 and R_2 returns are imputed from Figures 5 and 6 of the 2010 working paper version of BBK. The last column shows moments from CRSP value-weighted S&P 500 index. Shown are means, medians, standard deviations, and Sharpe ratios for the full sample (1996:2-2009:10) and the two subsamples (1996:2-2002:12 and 2003:1-2009:10), as well as AR(1) coefficients from the GARCH(1,1) model, and parameter estimates and R^2 values from the market and Fama-French (FF) 3-factor models for the full sample.

Table 5. HORIZON EFFECTS IN SHORT-TERM ASSET RETURNS

	Rescaled Average Returns					p-values			
				$R_i - SP$		$P(R_i \leq SP)$		$P(R_i = SP)$	
Horizon	R_1	R_2	SP	R_1	R_2	R_1	R_2	R_1	R_2
A. Full Sample									
1 month	1.15	1.13	0.58	0.57	0.55	0.15	0.19	0.30	0.38
3 months	1.02	0.92	0.59	0.43	0.32	0.20	0.25	0.40	0.50
6 months	1.02	0.88	0.58	0.44	0.30	0.21	0.25	0.42	0.50
12 months	0.92	0.72	0.53	0.39	0.19	0.25	0.35	0.50	0.70
165 months	0.84	0.66	0.47	0.37	0.19	-	-	-	-
1 month log	0.84	0.66	0.47	0.37	0.19	0.24	0.38	0.49	0.75
B. First Half									
1 month	1.58	1.40	0.68	0.90	0.72	0.18	0.26	0.36	0.52
3 months	1.31	0.88	0.70	0.61	0.18	0.26	0.42	0.52	0.84
6 months	1.28	0.81	0.64	0.63	0.16	0.27	0.42	0.54	0.84
12 months	1.18	0.65	0.70	0.49	-0.05	0.33	0.52	0.66	0.96
1 month log	1.09	0.66	0.55	0.54	0.11	0.29	0.46	0.57	0.92
C. Second Half									
1 month	0.71	0.85	0.47	0.23	0.38	0.31	0.24	0.62	0.47
3 months	0.74	0.95	0.49	0.25	0.46	0.25	0.11	0.50	0.22
6 months	0.76	0.95	0.51	0.25	0.44	0.19	0.05	0.38	0.10
12 months	0.65	0.78	0.36	0.30	0.43	0.13	0.05	0.26	0.10
1 month log	0.58	0.65	0.38	0.20	0.27	0.33	0.29	0.66	0.59

Notes: This table reports average returns (in percent) of the two short-term assets (R_1 and R_2) and the S&P 500 index (SP), calculated as $(\mathbb{E}[R_{i,t+1} \cdots R_{i,t+n}])^{1/n}$, where $R_{i,t}$ is the asset gross return in month t and n is the compounding horizon. Overlapping windows are used. Also reported are the differences in average returns of the short-term assets and the index and the associated p-values (based on a one-tailed test) computed using Newey-West (1987) methodology with n lags. Full sample covers 1996:2-2009:10, first half is 1996:2-2002:12, and second half is 2003:1-2009:10.

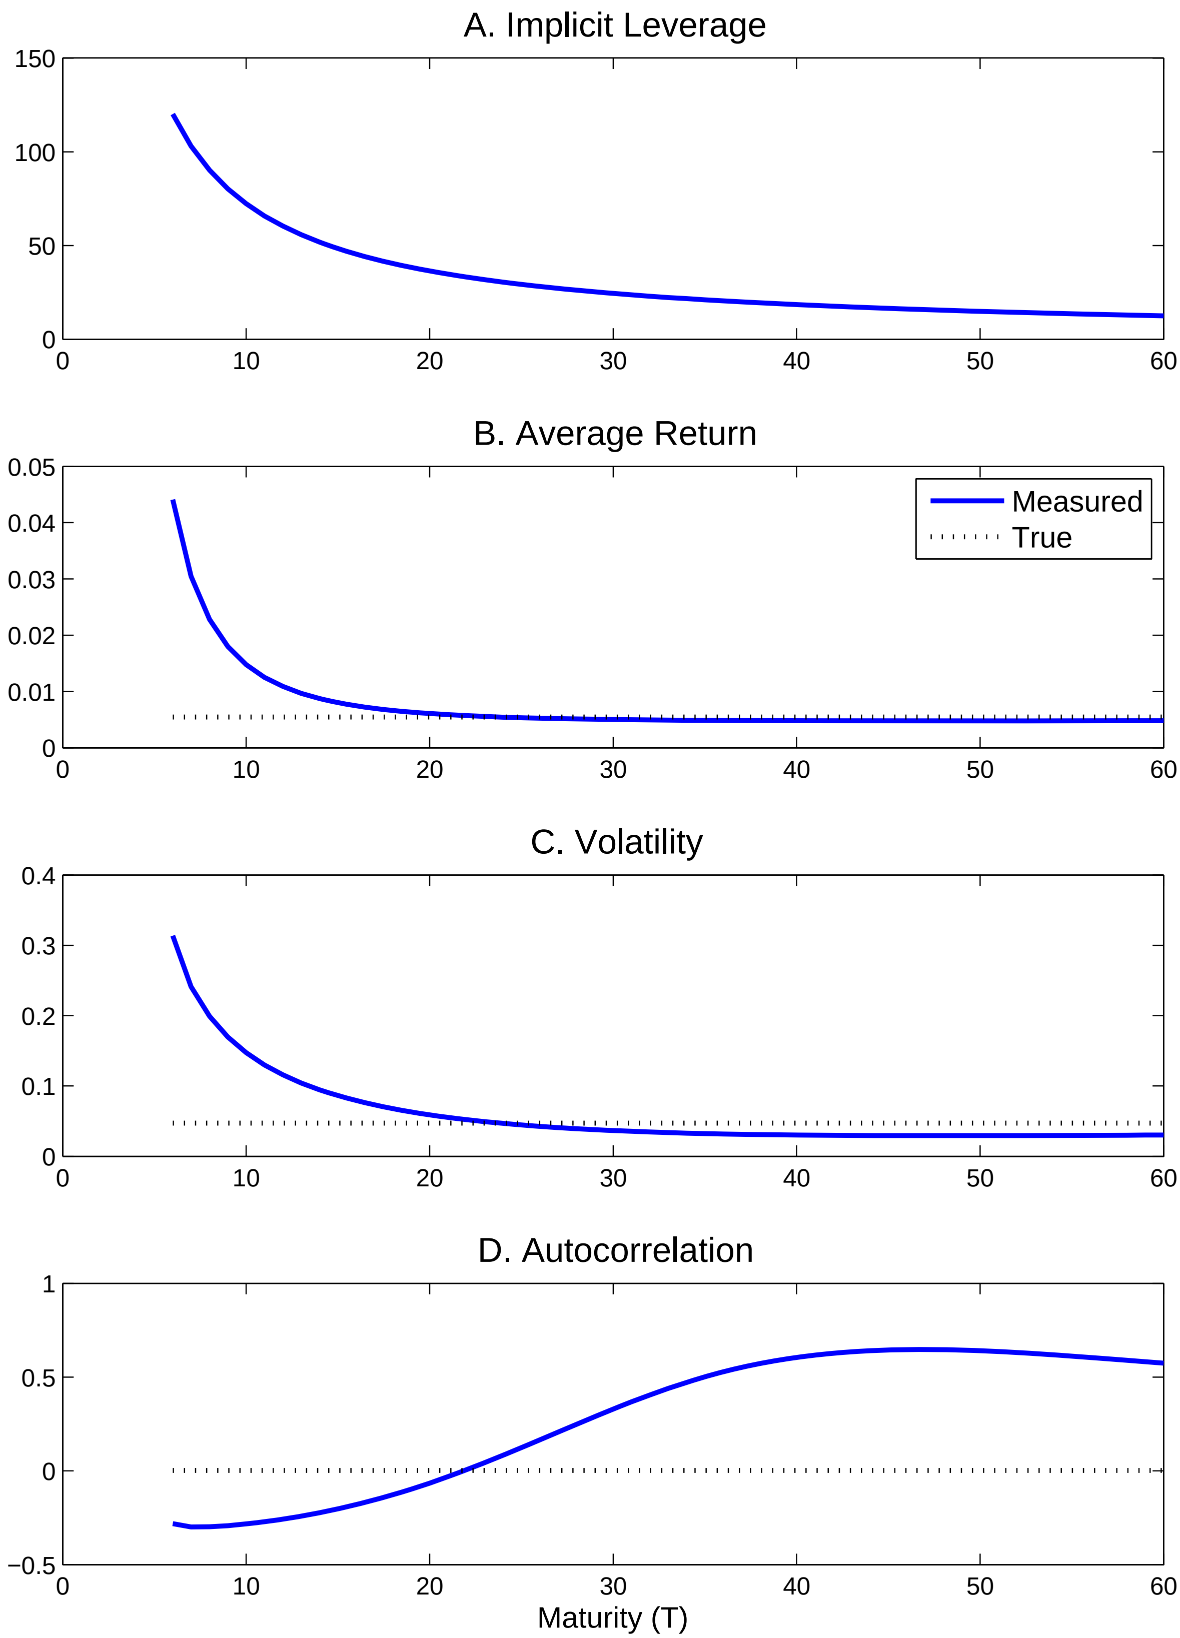

Figure 1. The Effects of Implicit Leverage on Dividend Strip Returns: Case I, Asynchronous Price Adjustment. This figure plots the average returns, volatility, and autocorrelation of the short-term asset return for different maturities T of the short-term asset. We use the base calibration of Table 3, Panel A, column (i), in which asynchronous prices cause small persistence in index returns.

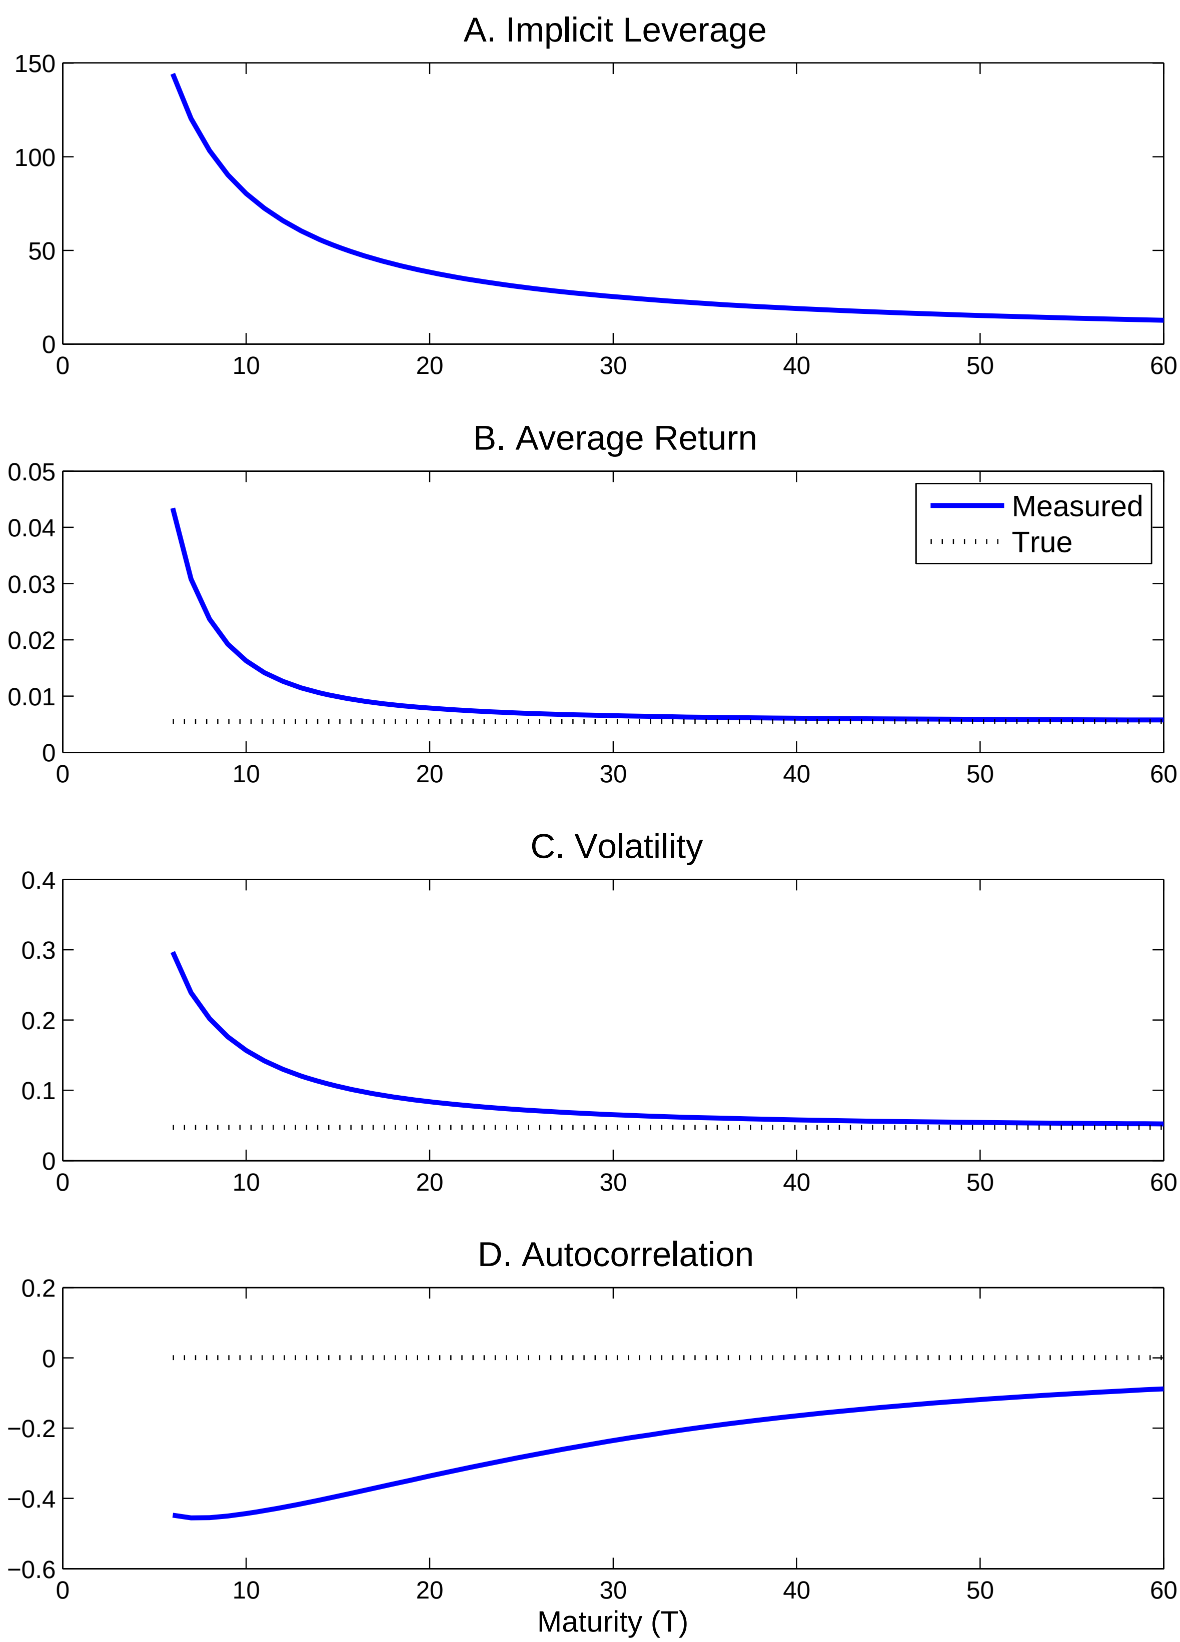

Figure 2. The Effects of Implicit Leverage on Dividend Strip Returns: Case II, Measurement Error. This figure plots the average returns, volatility, and autocorrelation of the short-term asset return for different maturities T of the short-term asset. We use the base calibration of Table 3, Panel B, column (i), in which measurement error impacts futures prices. We set the maturity T_1 of the long position to one month and the maturity T_2 of the short position varies along the horizontal axis of the figure.

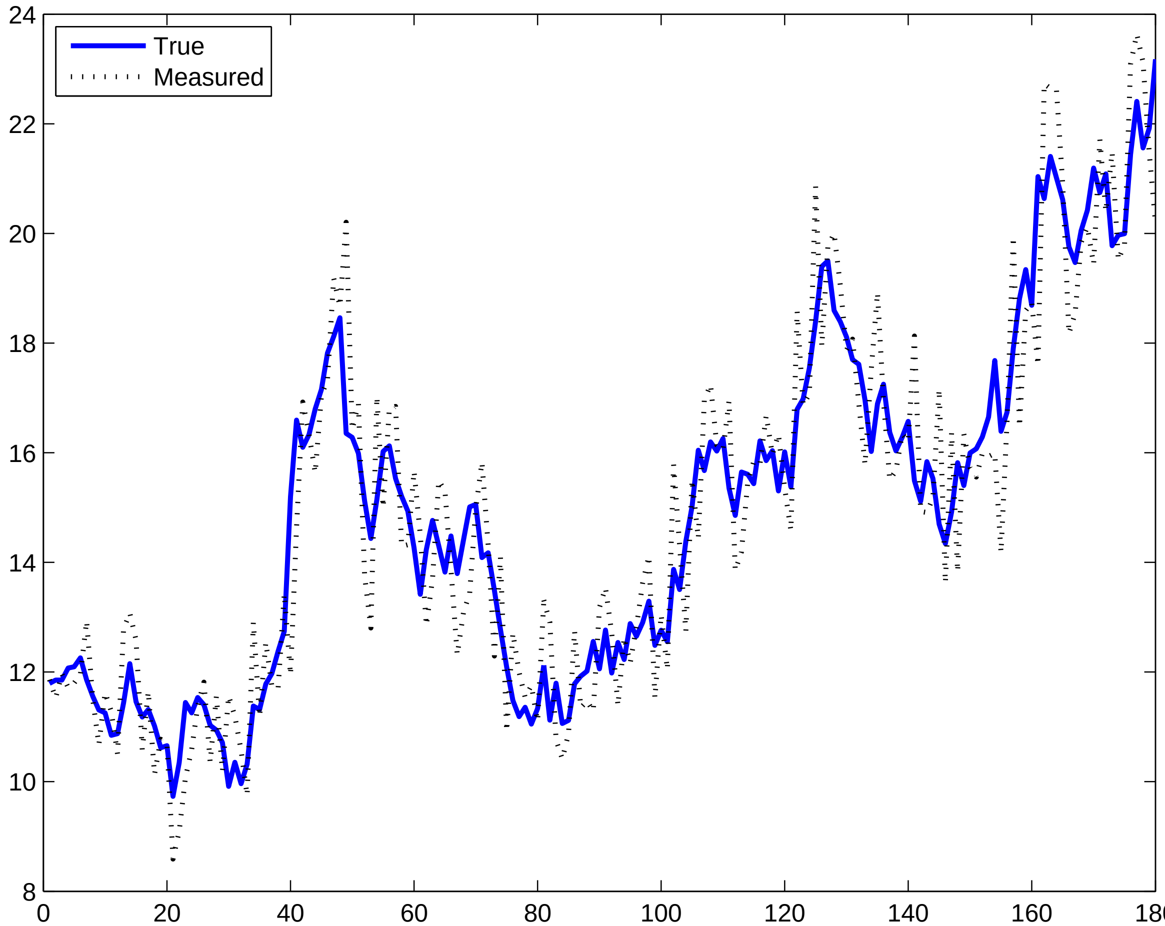

Figure 3. Cointegration between True and Measured Dividend Strip Prices. This figure plots simulated true and measured dividend strip prices using 180 months of data drawn from the calibration given in Table 3, Panel C, column (ii). The correlation of the two price series is high (approximately 94%), but the correlation of returns is much lower (20%). The unconditional average of the measured return series is approximately twice as high as the unconditional average of the true return series. (See Table 3.)

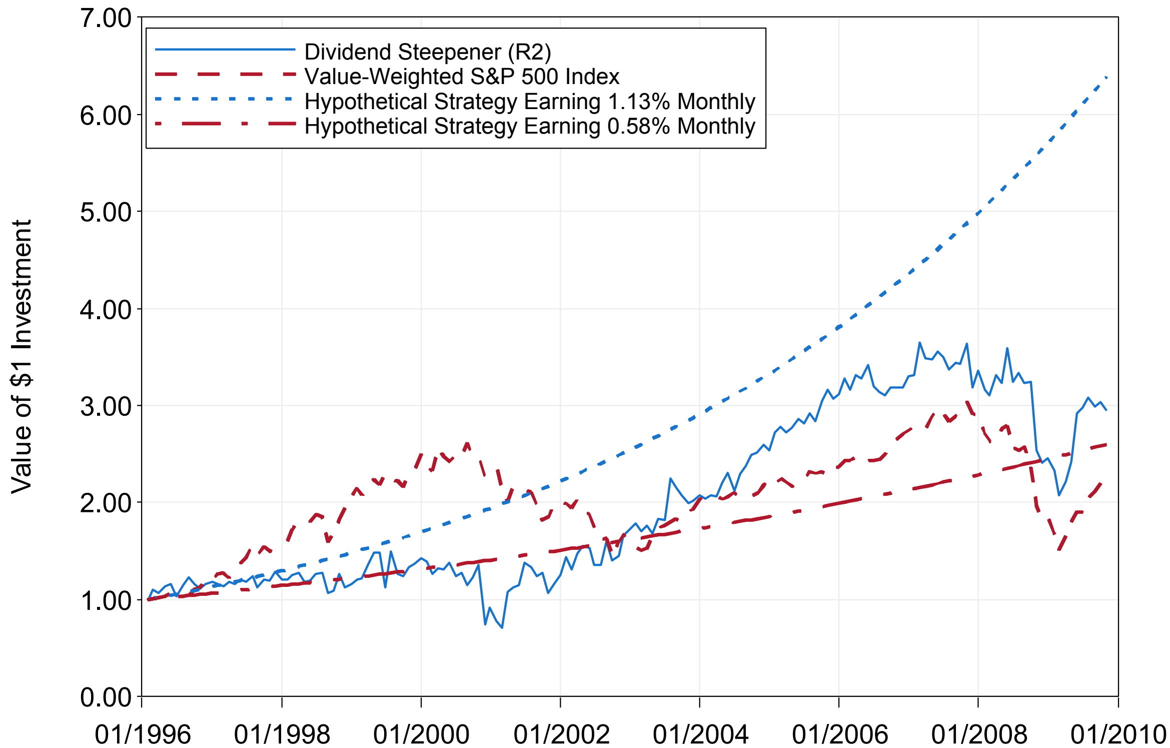

Figure 4. Comparison of Investment Performance. This figure plots the value of a $1 investment in the dividend steepener strategy (R_2), the value-weighted S&P 500 index, and two hypothetical strategies whose returns each month equal to the average returns of the dividend steepener strategy (1.13%) and the S&P 500 index (0.58%).