University of Kent

Kent Academic Repository

Quaye, Enoch, Tunaru, Radu and Voukelatos, Nikolaos (2025) *MIDAS and Dividend Growth Predictability: Revisiting the Excess Volatility Puzzle.* Journal of Financial Research, 48 (1). pp. 295-319. ISSN 0270-2592.

Downloaded from

https://kar.kent.ac.uk/105702/ The University of Kent's Academic Repository KAR

The version of record is available from

https://doi.org/10.1111/jfir.12403

This document version

Author's Accepted Manuscript

DOI for this version

Licence for this version

UNSPECIFIED

Additional information

Versions of research works

Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in ***Title of Journal*** , Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries

If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our Take Down policy (available from https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies).

MIDAS and Dividend Growth Predictability: Revisiting the Excess Volatility Puzzle

Enoch Quaye[*], Radu Tunaru[†] and Nikolaos Voukelatos[‡]

Abstract

We examine dividend growth predictability and the excess volatility puzzle across a large sample of international equity markets, using a mixed frequency data sampling (MIDAS) regression approach. We find that accounting for dividend seasonality under the MIDAS framework significantly improves dividend growth predictability, compared to simple regressions with annually aggregated data. Moreover, variance bounds tests that allow for non-stationary dividends consistently fail to reject the hypothesis of market efficiency across all countries. Our findings suggest that the common rejection of market efficiency in the previous literature is most likely driven by the annual aggregation of dividend data as well as by the assumption of stationary dividends.

Keywords: dividend growth predictability; seasonality; excess volatility puzzle; mixed frequency data
JEL Classifications: G12; G14; G17

1 Introduction

There are very few topics in finance and economics that have attracted as much attention as the question of what determines stock prices and, in particular, whether or not stock prices reflect the fundamental value of the underlying firms. The debate around this question has been running the full spectrum from Keynes' views about stock markets operating

[*]School of Accounting and Finance, University of Bristol. E-mail: Enoch.Quaye@bristol.ac.uk; Address: University of Bristol, Bristol BS8 1TU, UK

[†]University of Sussex. E-mail: R.Tunaru@sussex.ac.uk; Tel:+44 (0)1273 876920; Address: Room 206, Jubillee Building, University of Sussex, Brighton BN1 9RH, UK

[‡]Corresponding author; Kent Business School, University of Kent. E-mail: N.Voukelatos@kent.ac.uk; Tel:+44 (0)1227 827705; Address: Room 333, Sibson, Kent Business School, Canterbury CT2 7FS, UK

as casinos for the lucky to the Efficient Market Hypothesis of Fama and Samuelson. One specific approach that has been adopted to evaluate the rationality of stock markets refers to the variance bounds tests that were first introduced by Shiller (1981) and LeRoy and Porter (1981). Under this framework, the volatility of realized stock prices theoretically should not exceed the volatility of rational ex-post prices. However, subsequent empirical studies have commonly found that the variance bound is violated in the US market, with this finding often referred to as the excess volatility puzzle. This apparent rejection of market efficiency, though, has been based on comparisons against the volatility of expected prices that are, in turn, obtained using specific estimates of expected dividend growth. As a result, this type of market efficiency tests are predicated on being able to accurately forecast future dividend growth.

The literature on dividend growth predictability is characterized by a lack of consensus. Following Campbell and Shiller (1988), most studies have explored the extent to which dividend growth can be predicted by using the dividend-price ratio, failing to obtain consistent results across different countries and periods (Ang and Bekaert, 2007; Cochrane, 2008; van Binsbergen and Koijen, 2010). Therefore, recent studies have suggested that dividend growth predictability is weak at best, mainly as a result of dividend smoothing by large firms (Rangvid et al., 2014). However, Asimakopoulos et al. (2017) argue that the reported lack of dividend growth predictability is driven by the use of annually-aggregated dividends, with time aggregation discarding potentially useful information about future dividend growth.

In this paper, we adopt a new approach for evaluating market efficiency by accounting for intra-year seasonality within a mixed frequency data sampling (MIDAS) regression setting. This methodology is motivated by the well-known seasonality in dividend payments, with the peak of activity around the April-June period. While previous studies have used annually-aggregated dividends to compute expected dividend growth, our use of the MIDAS approach allows the variance bounds tests to take into account the important information contained in monthly dividends.

Our paper also contributes to the literature on dividend growth predictability and excess volatility by shifting the focus from the US to a large sample of international stock markets. Previous studies have tended to focus almost exclusively on the US stock market. While this emphasis on the US is certainly understandable, we believe that expanding the analysis to an international context can provide valuable new insights into the nature of dividend growth predictability. Moreover, revisting the excess volatility puzzle in an international context can allow us to reach more meaningful conclusions about the efficiency of equity markets at different stages of development. To this end, we examine a mix of 50 sample countries that are substantially disperse, in terms of both geographical location as well as state of economic development. To the best of our knowledge, this is the first study of dividend growth predictability and the excess volatility puzzle to examine a large sample of international markets and the first one using MIDAS regression in this context.

Under this new modelling approach, we provide evidence that stock markets with a short history tend to be inefficient but, as the sample size increases, the volatility ratio that is the basis of variance bound tests approaches the threshold level of one that is consistent with market efficiency.

One of our main findings is that using time-disaggregated (monthly) dividends to compute the dividend-price ratio results in a markedly higher predictive power over future dividend growth, compared to that obtained from annually-aggregated dividends (consistent with Asimakopoulos et al., 2017). Consequently, accounting for dividend seasonality under the MIDAS framework provides significantly more meaningful measures of expected dividend growth and, by extension, more meaningful measures of ex-post rational stock prices that are used in the variance bounds tests. We show that the results of variance bounds tests are highly sensitive to the specific approach used for obtaining expected dividend growth, highlighting the importance of accounting for intra-year dividend seasonality.

We also report evidence of varying performance of different weighting schemes in the MIDAS regressions. In particular, we consider a set of different MIDAS weighting schemes, and we base our selection of the optimal scheme for each country on an out-of-sample forecasting exercise. We find that in the majority of countries, the highest predictive power of the dividend-price ratio over subsequent dividend growth is produced when the weights are given as an Almon lag polynomial of order P. Moreover, we show that dividends that are paid out towards the end of the year take substantially higher weights in the MIDAS estimations compared to dividends that are paid out earlier in the year, confirming the importance of taking into account the role of different dividend lags under the MIDAS approach.

When we perform the first generation variance bounds test proposed by Shiller (1981), we find strong evidence of excess volatility in our sample of international stock markets, in line with the excess volatility puzzle in the US documented by Shiller (1981) and LeRoy and Porter (1981). More specifically, the volatility of realized prices exceeds in various degrees the volatility of ex-post rational prices computed from the subsequent dividends in all sample countries. Furthermore, in the majority of cases, the ratio of realized price volatility over the volatility of ex-post rational prices exceeds the value of 5:1 that was reported in Shiller (1981) for the US market, suggesting that deviations from the hypothesis of market rationality could potentially be even more pronounced in an international context.

However, an important feature of the first generation bounds test that has been called into question is the assumption of stationary dividends. For instance, Engel (2005) derives a bounds test on the variance of first differences in prices, assuming that dividends can follow either a stationary or a unit-root process and that the arithmetic price-change variance is a monotonically decreasing function of investors' information about future dividends. Under this set of assumptions, Engel (2005) shows that the excess variance inequality would in fact be reversed.[1]

We find that the assumption of stationary dividends is consistently violated in our international sample, and we also report evidence of small sample bias in the first generation test results. Therefore, we base our conclusions about market efficiency on the second generation bounds test proposed by Engel (2005). When we allow for the observed deviation of dividends from stationarity, the variance bounds tests fail to detect excess volatility in any of the 50 sample countries, consistent with the hypothesis of market efficiency. Importantly, we report that the variance of realized stock price changes is significantly closer to that of expected price changes when the latter is obtained via MIDAS regressions, compared to regressions with annually-aggregated dividends (as in Cochrane, 2008). In this sense, our empirical findings provide strong evidence that bounds tests' rejection of market efficiency in previous studies is most likely driven by dividend non-stationarity and the discarding of information from the annual aggregation of dividends, as opposed to reflecting genuine market inefficiencies.

The remaining of this study is organized as follows. Section 2 discusses the data used in the empirical analysis and the construction of the main variables of interest. Section 3 presents the MIDAS framework and the methodology for selecting the optimal weighting scheme per country. Section 4 discusses the variance bounds tests of market efficiency. Section 5 presents the empirical results, while Section 6 concludes.

2 Data Sources and Main Variables

We obtain international data from Global Financial Data (GFD). We cover 50 countries and the data includes monthly observations for the nominal stock index price, stock index dividend, the Consumer Price Index (CPI), and the risk-free rate for each country in the sample.[2] The overall sample period runs from January 1840 to December 2018, for a total of 179 years (2,148 months). However, given that coverage in GFD begins at different times for different countries, the number of available observations varies across the sample countries. For example, data for the French equity market is available from 1840 while, at the other end of the spectrum, coverage for Bulgaria only begins in 2001.

The MIDAS estimation involves regressing the dividend growth (at an annual frequency) against the lagged dividend-price ratio (at a monthly frequency). The monthly quoted variables $P^m_{t,k}$ and $D^m_{t,k}$ denote the price levels and dividend levels, respectively,

[1] More recently, Lansing and LeRoy (2014) demonstrate that the relationship between log-return variance and investor information can be non-monotonic, depending on the specific nature of investors' risk aversion. In addition, Lansing (2016) modifies the framework of Engel (2005) regarding investors' risk preferences and information about future dividends to provide an alternative bounds test on the variance of first differences in prices.

[2] See the first column in Table 4 for the full list of countries in our paper.

observed in year t and month k. The variables $p^m_{t,k}$, $d^m_{t,k}$, and $y^m_{t,k}$ denote the corresponding logarithmic series of prices, dividends, and the dividend-price ratio, respectively. For each country, we obtain annual observations by aggregating monthly dividends, with the time-series of annual dividends D_t being given as the sum of the 12 monthly dividends that were paid during a particular year. Finally, we compute real annual prices and dividends by deflating the nominal time-series by the respective CPI.[3]

The 50 sample countries have experienced positive mean dividend growth during the sample period, with the only exception of Egypt. Despite its almost universally positive sign, mean dividend growth varies substantially across different countries, ranging from a minimum of 0.013 (Japan) to a maximum of 0.465 (Argentina). Interestingly, countries with shorter available data series generally have a higher mean and standard deviation of dividend growth compared to countries with more available data. For example, the highest standard deviation of dividend growth is exhibited by Bulgaria (0.994) which only has 17 years of data in our sample. Similarly, the mean log return also seems to be negatively related to the number of available observations across countries, and it remains predominantly positive across sample countries (with the exception of Portugal).[4]

We perform a number of stationarity and unit-root tests on dividends and prices. More specifically, we run the Augmented Dickey-Fuller (ADF) test, the Phillips-Perron (PP) test, and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test. Table 1 reports the number of sample countries that reject the null hypothesis in each test.[5] Our results are consistent with those reported in previous studies about dividends deviating from stationarity, with the KPSS test rejecting the null hypothesis of dividends following an $I(0)$ process across all 50 countries at the 10% level, while the null of stationarity is rejected at the 5% level in 34 countries. Moreover, the results of the ADF test suggest that dividends generally follow an $I(1)$ process, with the null of a unit root being rejected at the 5% level in only 4 countries (Chile, Malaysia, Italy, and Canada). The PP test results are very similar, with the null of a unit root being rejected at the 5% level in only 5 out of 50 countries (Taiwan, New Zealand, Italy, Australia, and Canada).

[Table 1 about here.]

[3] We follow the approach in Shiller (1981) to deflate nominal variables. Thus, we use the mean annual CPI to deflate monthly variables for observations before 1900 and monthly CPI to deflate observations after 1900.

[4] Unreported plots of the time-evolution of index price levels and the corresponding dividends depict a clear upward trend in both series across the 50 sample countries, with a noticeable degree of co-movement between the two types of time-series. Descriptive statistics on the time-series of index returns, dividend growth and the dividend yield in each country are reported in Table A1 in the Internet Appendix, in order to conserve space.

[5] The results of the stationarity and unit root tests are reported in Table B1 in the Internet Appendix, to conserve space.

We also perform a set of pool unit root tests on the original and de-trended time-series of prices and dividends. More specifically, Table 2 reports the results of the Im-Persaran-Shin, Augmented Dickey-Fuller, and Phillips-Perron unit root tests using pooled data across all 50 sample countries. Overall, the results reject the hypothesis of a unit root in the pooled dataset.

[Table 2 about here.]

3 MIDAS and Expected Dividend Growth

3.1 Expected Dividend Growth

Since rationally expected future dividends are unobservable by nature, it is important to use an appropriate rate of expected dividend growth when we compute the rational stock price as the discounted value of future dividends. In the seminal study by Shiller (1981), the expected growth rate for stock prices (and, by extension, for dividends) is given by the trend factor from regressing stock prices against time

$$p_t = \alpha + \beta t + \epsilon_t \tag{1}$$

However, the validity of the regression in equation (1) hinges on the stationarity of prices. In our sample, the ADF, PP and KPSS tests indicate that prices deviate significantly from stationarity across virtually all 50 countries, suggesting that a growth factor that is based on regressing prices against a deterministic trend is unlikely to be appropriate.[6] More recent studies on dividend predictability suggest that alternative measures of expected dividend growth might be better suited to forecast future dividends and, hence, could provide a more efficient ex-post rational price. To this end, we explore the predictability of dividend growth and the extent to which it can be forecasted using the dividend-price ratio (see Cochrane, 2008; Asimakopoulos et al., 2017; Golez and Koudijs, 2018). For each country, we run separately a time-series regression of dividend growth on the previous period's dividend-price ratio as follows

$$\Delta g_{t+1} = \beta_{0,g} + \beta_{1,g} y_t + \epsilon_{g,t+1} \tag{2}$$

where the subscript t denotes time at an annual frequency. Following Cochrane (2008), the standard errors are GMM-corrected for heteroscedasticity. Similarly to Ang and Bekaert

[6] In the interest of comparability with some of the earlier literature, we repeat all the subsequent empirical analysis using the de-trending factor obtained from Shiller (1981) regressions of prices against time, ignoring the issue of non-stationarity of prices. The results are reported in the Internet Appendix, to conserve space.

(2007), we begin by estimating these time-series regressions using aggregated annual data, ignoring seasonality issues. By adopting an approach based on more recent evidence on dividend predictability, we aim to obtain an improved de-trending factor $\lambda = exp(\beta)$, with the slope from equation (2) replacing the slope that would have been obtained from regressing stock prices against time.

An important concern at this point is that the aggregation procedure involved in computing annual prices and dividends to use in equation (2) will lead to some loss of within-year information. We address this concern by employing a mixed-frequency regression framework (MIDAS), where the lower-frequency (annual) dependent variable Δg_{t+1} is regressed against a higher-frequency (monthly) independent variable y_t^m.

3.2 The MIDAS Framework

The MIDAS approach was introduced in Ghysels et al. (2005) and Ghysels et al. (2006) as a framework to estimate regression specifications where the dependent and the independent variables are sampled at different frequencies. Subsequent empirical studies have used the MIDAS approach in a variety of applications where the dependent variable is typically quoted at a lower frequency compared to the independent variables. For example, Forsberg and Ghysels (2007) estimate a MIDAS regression specification in the context of forecasting index volatility at longer horizons using absolute daily returns, while Clements and Galvão (2008), Bai et al. (2013), and Gagliardini et al. (2017) use monthly macroeconomic and financial indicators to forecast quarterly GDP growth. Our use of the MIDAS approach is more closely related to Asimakopoulos et al. (2017) who explore the predictive ability of the monthly dividend-price ratio over subsequent annual dividend growth.

In our study, the MIDAS framework is applied to estimate the effect of the higher frequency data of the log dividend yield y_t^m (monthly, i.e. $m = 12$) on the lower frequency data of the log dividend growth rate Δg_{t+1} (annual). The applied regression model is

$$\Delta g_{t+1} = \widehat{\beta}_{0,g} + \widehat{\beta}_{1,g} B(L^{1/m}; \boldsymbol{\theta}) y_t^m + \widehat{\epsilon}_{g,t+1} \tag{3}$$

for $t = 1, \ldots, T$, where $L^{1/m}$ denotes the lag operator of the log dividend yield data. The term

$$B(L^{1/m}; \boldsymbol{\theta}) = \sum_{k=0}^{K-1} \omega_k(\boldsymbol{\theta}) L^{k/m} \tag{4}$$

denotes a known polynomial function of $L^{1/m}$ whose coefficients depend on a small dimensional vector of parameters θ, while $L^{k/m}$ is the lag operator of y_t^m for k/m periods. The maximum length of the polynomial function is $K - 1$. The overall impact of the lagged

y_t^m on Δg_{t+1} is given by the coefficient $\widehat{\beta}_{1,g}$, which can be obtained by normalizing the weights $\omega_k(\boldsymbol{\theta})$ so they sum up to one.[7] The general form equation (3) can be rewritten as

$$\Delta g_{t+1} = \widehat{\beta}_{0,g} + \widehat{\beta}_{1,g}\omega\widehat{y}_t^m + \widehat{\epsilon}_{g,t+1} \tag{5}$$

where $\omega\widehat{y}_t^m = \sum_{k=0}^{11} \omega_k y_{t,k}^m$. Equation (5) can be thought of as a projection of the annual dividend growth variable Δg_{t+1} onto the monthly log dividend yield variable y_t^m using up to $K-1$ monthly lags (i.e. 11 lags in our case).

3.3 MIDAS Weighting Schemes

The specific shape of the weighting scheme ω will depend on the chosen specification for the polynomial function. On this issue, Ghysels et al. (2007) discuss a number of alternative weighting schemes. In this study, we consider four alternative specifications for the polynomial function, each of which will result in a different weighting scheme and, by extension, a different effect of y_t^m on Δg_{t+1}. Our pool of candidate weighting schemes consists of a polynomial with a step-function, an exponential Almon lag polynomial, a normalized beta lag polynomial, and an Almon lag polynomial of order P.

The first specification is a polynomial with a step-function (see also Forsberg and Ghysels, 2007). Under this alternative, a regressor X_t can be expressed as the partial sum of the higher frequency x^m, so that

$$X_t(K, m) = \sum_{j=1}^{K} x_{t-\frac{j}{m}}^m. \tag{6}$$

Then, the MIDAS regression with M steps can be estimated as a simple Ordinary Least Squares (OLS) regression of the lower-frequency dependent variable against the regressor $X_t(K_i, m)$, where $K_1 < \cdots < K_M$. The impact of x^m on the dependent variable can be measured by the sum of all coefficients in the OLS regression (i.e. $\sum_{i=1}^{M} \beta_i$), since it appears in all the partial sums. A more detailed discussion of MIDAS with step-functions can be found in Forsberg and Ghysels (2007).

The second specification is an exponential Almon lag polynomial, similarly to the one used in Asimakopoulos et al. (2017). Ghysels et al. (2007) argue that this scheme can be thought of as the most general weighting scheme, as it has the most flexible shape. In its unrestricted version, the exponential Almon lag polynomial is fully determined by its two parameters θ_1 and θ_2, with the corresponding weights computed as

[7] See Ghysels et al. (2006) for a detailed discussion on the benefits of weight normalization in MIDAS models.

$$\omega_k(\theta_1, \theta_2) = \frac{e^{\theta_1 k + \theta_2 k^2}}{\sum_{k=1}^{K} e^{\theta_1 k + \theta_2 k^2}} \tag{7}$$

The third specification is a normalized beta lag polynomial (see also Asgharian et al., 2013). Based on the beta function, this polynomial is fully determined by the three parameters θ_1, θ_2 and θ_3. The beta function is very flexible, as it allows for weights that can take a variety of different shapes. For instance, Ghysels et al. (2007) show that larger values of θ_2 result in faster declining weights, with the rate of weight decline essentially determining how many lags will be included in the MIDAS specification. The weights under the normalized beta lag polynomial can be computed as

$$\omega_k(\theta_1, \theta_2, \theta_3) = \frac{h_k^{\theta_1 - 1}(1 - h_k)^{\theta_2 - 1}}{\sum_{k=1}^{K} h_k^{\theta_1 - 1}(1 - h_k)^{\theta_2 - 1}} + \theta_3 \tag{8}$$

where $h_k = (k-1)/(K-1)$. In addition to its unrestricted version above, we also consider a restricted version where the last lag is set to zero, i.e. $\omega_k(\theta_1, \theta_2, 0)$.

The final specification is the non-normalized Almon lag polynomial of order P, first introduced in Almon (1965). The corresponding weights for each lag k are computed as

$$\omega_k(\theta_0, \dots, \theta_P) = \sum_{p=0}^{P} \theta_p k^p \tag{9}$$

The weights in (9) are obtained via a non-linear least squares estimation, where the optimal lag order is selected using the AIC/BIC of the least squares estimation (for more details, see Ghysels et al., 2007). This approach assumes that the successive weights lie on a polynomial, estimating a few points on the curve as regression coefficients and then using polynomial interpolation to interpolate between them for the remaining points.

3.4 Selection of Weighting Scheme

Instead of selecting a single weighting scheme universally across all countries, we adopt a more flexible approach where we select the optimal weighting scheme separately for each country, from the pool of four schemes presented above. This country-by-country selection of the optimal weighting scheme is based on an out-of-sample evaluation of the forecasts produced by the candidate schemes (similar to Andreou et al., 2013).

For each sample country, we begin by constructing an estimation period that starts at the beginning of the respective data series and ends 2 years before the end of the available data. We use this estimation period to estimate the MIDAS model under each candidate scheme for a given country. Then, we use the remaining 2-year period in order to produce out-of-sample forecasts of dividend growth under each candidate scheme. Finally, we select the optimal scheme, separately for each sample country, as the one that produces the lowest out-of-sample Root Mean Squared Error (RMSE).

Table 3 reports descriptive statistics of the RMSE for each weighting scheme across the 50 sample countries. On average, the Almon lag polynomial of order P is found to produce the lowest RMSE across the 50 countries, with a mean RMSE of 0.070, while the Exponential Almon polynomial produces the largest mean RMSE (0.105). However, the polymomial with a step-function results in the lowest median RMSE across the 50 countries (at a median level of 0.062). More importantly, we find that the Almon P polynomial is the most often selected weighting scheme, resulting in the lowest RMSE among the candidate schemes for 20 out of 50 countries. The second most often selected scheme is the normalized beta polynomial, either unconstrained (selected in 8 cases) or in terms of the special case where the last lag is constrained to zero (selected in 6 cases).[8]

[Table 3 about here.]

Regarding the relative importance of different lags, the MIDAS results indicate that the dividends paid towards the end of the year have a substantially greater role when predicting the annual dividend growth, compared to dividends that are paid out earlier in the year. For instance, similarly to Asimakopoulos et al. (2017), we find that dividends that are paid out in December consistently take the highest weight, while dividends that are paid out in the summer months tend to take the lowest weights. Figure 1 illustrates this stylized fact by plotting the weighting schemes across four sample countries (US, UK, France, and Spain) using the Almon lag polynomial of order P. This highlights the merits of MIDAS in terms of accounting for dividend seasonality and allowing different lags of the dividend-price ratio to take different weights when predicting dividend growth.

[Figure 1 about here.]

Overall, one significant advantage of the MIDAS approach is its flexibility in that it does not impose any particular assumptions about the effect of different lags. Instead, the optimal weighting scheme is driven entirely by the data. Asimakopoulos et al. (2017) further highlight that the use of non-linear lag polynomials under a MIDAS regression results in a more parsimonious estimation with a lower sensitivity to specification errors, compared to the alternatives of state-space models or mixed-frequency vector autoregression

[8]Interestingly, we find that every MIDAS weighting scheme tends to produce lower RMSEs for countries with longer series of available data. For example, the RMSE of the unconstrained beta polynomial is 0.49 in the case of Bulgaria (with only 17 years of available data) and 0.03 for Canada (with 187 years of available data). Table C1 in the Internet Appendix reports the RMSE per weighting scheme for each country.

(VAR) models. Finally, the MIDAS approach avoids the issue of parameter proliferation, which could have more pronounced consequences in some of our sample countries with relatively small datasets.

[Figure 2 about here.]

4 Variance Bounds Tests

The log dividend yield can be expressed as the difference between expected stock returns and the expected dividend growth plus a constant (Campbell and Shiller, 1988), given by

$$y_t = \alpha + E_t[\sum_{j=0}^{\infty} \rho^{j-1} r_{t+1+j}] - E_t[\sum_{j=0}^{\infty} \rho^{j-1} \Delta g_{t+1+j}] \tag{10}$$

where ρ denotes an autoregressive parameter. We follow Cochrane (2008) and estimate ρ via a Vector Autoregressive (VAR) specification based on the changes in dividend growth Δg_t, stock returns r_t, and the dividend yield y_t. This specification can serve as a reasonable starting point which allows us to obtain ex-post rational prices by using the dividend yield and expected dividend growth.

As discussed earlier, while the dividend yield is readily observable at time t, we still need a meaningful measure of expected dividend growth based on information available at t. Hence, in the empirical analysis we compute the growth trend obtained via a regression of dividend growth against the lagged dividend-price ratio using annually-aggregated data (as in Cochrane, 2008) or mixed-frequency data (as in Asimakopoulos et al., 2017). More specifically, we obtain the long-run exponential growth rate that is used to de-trend the price and dividend series as $\lambda = e^{\beta}$, where β refers to the slope coefficient from the regression in (2). The real de-trended time-series of prices and dividends are, then, given by $\widehat{p}_t = \frac{p_t}{\lambda^{t-T}}$ and $\widehat{d}_t = \frac{d_t}{\lambda^{t+1-T}}$, respectively.

Finally, similarly to Shiller (1981), we compute the de-trended ex-post rational price $\widehat{p}_t^*$ recursively from the terminal date T using the equation

$$\widehat{p}_t^* = \bar{\gamma}(\widehat{p}_{t+1}^* + \widehat{d}_t) \tag{11}$$

where $\bar{\gamma} = \lambda(1+r)$ is a discount factor and r denotes the one-year risk-free rate of interest. In order to solve the recursive problem in (11), we set the terminal value $\widehat{p}_T^*$ of the ex-post rational price $\widehat{p}_t^*$ equal to the average de-trended realized price over the sample period, i.e $\widehat{p}_T^* = \frac{1}{T}\sum_{t=1}^{T} \widehat{p}_t$.

The efficient markets model implies that, since rational investors determine current stock prices by discounting future dividends, $\widehat{p}_t$ represents an optimal forecast of $\widehat{p}_t^*$, in other words $\widehat{p}_t = E_t[\widehat{p}_t^*]$. Then, it follows that the forecast error $u_t = \widehat{p}_t^* - \widehat{p}_t$ must be uncorrelated with the forecast itself. This means that $var(\widehat{p}_t^*) = var(\widehat{p}_t + u_t) = var(\widehat{p}_t) + var(u_t)$. Since variances are obviously non-negative, this relationship results in the following expression for the first generation bounds test of

$$\sigma(\widehat{p}_t) \leq \sigma(\widehat{p}_t^*) \tag{12}$$

Shiller (1981) and LeRoy and Porter (1981) were the first studies to document that this upper bound is violated in the US market, with the volatility of realized prices exceeding that of ex-post rational prices by a factor of 5. In addition to the upper bound of the variance of stock prices that is described in (12), the maximum value of the variance of changes in price for a given variance of dividends is given by

$$\sigma(\Delta\widehat{p}) \leq \frac{\sigma(\widehat{d})}{\sqrt{2r}} \tag{13}$$

The main intuition behind the proof of the inequality in (13) is that the variance of changes in price is larger when information about future dividends is revealed more smoothly across time, as opposed to future dividends being known either many years before or just before they are paid.[9] Finally, a slightly different version of the variance inequality can be written as

$$\sigma(\Delta\widehat{p} + \widehat{d}_{t-1} + r\widehat{p}_{t-1}) \leq \frac{\sigma(\widehat{d})}{\sqrt{2r}} \tag{14}$$

The theoretical variance inequalities in (12), (13) and (14) are based on the assumption that dividends follow a stationary process. However, the assumption of stationarity has since been called into question as being unlikely to characterize the true dividend process. For instance, Marsh and Merton (1986) argue that dividends are most likely non-stationary because of the general tendency of firms to smooth dividends over time, while Merton (1986) argues that the results in Shiller (1981) are more likely to reflect a rejection of the stationarity assumption than a rejection of market efficiency. Since we also find strong evidence of non-stationarity in the time-series of dividends across all 50 sample countries, we place greater emphasis on a set of second generation bounds tests that have been developed to account for, among other issues, the observed non-stationarity of dividends.

More specifically, West (1988) shows that, under relatively weak assumptions including potential non-stationarity in dividends, the variance of innovations in the stock price must

[9] Shiller (1981) uses a standard first-order autoregressive specification for future dividends to derive the maximum value for the variance of changes in price. We drop the time subscripts from the variance inequality in (13) because the unconditional co-variance between $\widehat{d}_t$ and any information variable will depend on k but not on t.

be lower than the variance of innovations in the corresponding dividend.[10] Engel (2005) extends the analysis by deriving a new variance bound on the first difference of stock prices, under the assumption that dividends are stationary or that they follow a unit-root process. Working with differences in prices, it is shown that the excess volatility inequality is reversed to

$$var(\Delta\widehat{p}_t) \geq var(\Delta\widehat{p}_t^*) \tag{15}$$

where $\Delta\widehat{p}_t = \widehat{p}_t - \widehat{p}_{t-1}$ and $\Delta\widehat{p}_t^* = \widehat{p}_t^* - \widehat{p}_{t-1}^*$.

5 Empirical Results

5.1 Dividend Predictability Regressions

Table 4 reports the results of dividend predictability regressions under the two alternative approaches to obtaining measures of expected dividend growth, namely regressions of changes in dividend growth against the lagged dividend-price ratio at an annual frequency (as in Cochrane, 2008) and MIDAS regressions of changes in annual dividend growth against the lagged monthly dividend-price ratio (as in Asimakopoulos et al., 2017). For comparability with the earlier literature, we also report the regression results from estimating the original regression of prices against time (as in Shiller, 1981).

[Table 4 about here.]

Estimating the standard Shiller-type regressions of log-price against time in (1) results in universally positive slopes across all sample countries, with all 50 coefficients being statistically significant at the 1% level. Moreover, the magnitude of these slopes is generally lower for developed countries with longer time-series relative to developing countries with shorter time-series. For instance, Japan has 117 years of available data and it is found to have the lowest slope ($\beta = 0.03$), compared to Romania which has only 19 years of data and the highest slope ($\beta = 0.64$). This relationship suggests that developed countries experience on average a lower stock price growth compared to developing countries.

Regressing changes in dividend growth against the lagged dividend-price ratio produces predominantly positive slope coefficients. More specifically, when we use annual dividend and price figures to estimate the regression specification in (2), 35 out of 50 countries are found to have a positive $\widehat{\beta}_{1,g}$, with 11 of these cases being statistically significant at the 5% level. By comparison, out of the 15 negative slopes, only 3 are found to be

[10] West (1988) computes innovations, i.e. unexpected changes, in the stock price and the corresponding dividend series as the residuals from fitting ARIMA (q, s, r) specifications on these series. In addition to relaxing the assumption of stationarity of dividends, the variance inequality proposed in West (1988) also attempts to address the issue of small sample bias.

statistically significant at the 5% level (namely for Austria, Finland, and Japan). When we account for intra-year seasonality by estimating the MIDAS specification in (3), the resulting slope coefficients of the dividend-price ratio are now universally positive, with 15 out of 50 positive slopes also being statistically significant at the 5% level.

In addition, Table 5 reports descriptives statistics for the R-square obtained from estimating each regression across the 50 sample countries. In terms of in-sample predictability, accounting for seasonality by using a mixed frequency regression setting seems to improve the measure of expected dividend growth considerably relative to using annually-aggregated values. In particular, the goodness-of-fit of the MIDAS regressions in (3) is substantially higher relative to that of the Cochrane-type regressions in (2), with a mean R-square of 55% in the former compared to only 4% in the latter. Furthermore, the R-square in the MIDAS regressions ranges from a minimum of 9% to a maximum of 96% across the 50 countries, while the respective range of the R-square in the Cochrane-type regressions is 0% - 20%. Importantly, the MIDAS specification results in a higher R-square compared to the one obtained from the Cochrane-type specification in all 50 sample countries. Based on this difference in predictive power, we would expect the de-trending factor $\lambda = e^{\beta}$ to be more accurate when β is proxied by the slopes from the MIDAS regressions compared to those from the Cochrane-type regressions.[11]

[Table 5 about here.]

5.2 First Generation Bounds Tests

In order to get a preliminary idea about the magnitude of the excess volatility puzzle across different markets, we begin by plotting in Figure 3 the historical evolution of the de-trended stock price $\widehat{p}_t$ and the corresponding ex-post rational price $\widehat{p}_t^*$ for France, Spain, the US and the UK. The resulting plots are consistent with the hypothesis of excess volatility in the time-series of realized prices compared to that of dividend-based expected prices, as evidenced by $\widehat{p}_t^*$ consistently plotting as a much smoother and more stable series compared to that of its respective $\widehat{p}_t$. In this sense, Figure 3 provides some preliminary evidence against the rational expectations hypothesis in an international context, consistent with the respective figures for the US in Shiller (1981), which depict the time-series of ex-post rational prices as considerably smoother than the time-series of realized prices.

[Figure 3 about here.]

[11]Interestingly, regressing stock prices against time (as in Shiller, 1981) also results in a very high in-sample fit, with a mean R-square of 53%. Nevertheless, these R-square values are not directly comparable to those obtained when estimating the predictive regressions in (2) and (3), since they refer to regressions with different dependent variables. Moreover, the validity of the Shiller-type regressions hinges on the stationarity of prices, which we find to be violated to various degrees across all sample countries.

Figure 4 plots the ratio $\theta = \frac{\sigma(\widehat{p}_t)}{\sigma(\widehat{p}_t^*)}$ of the volatility of the realized price over the volatility of the ex-post rational price, which refers to the first variance inequality in (12), across the 50 sample countries. In the interest of comparability with the earlier literature, rational ex-post prices for Figure 4 are based on expected dividend growth that has been obtained via regressions of prices against time, as in Shiller (1981).[12]

The first thing to notice is that $\sigma(\widehat{p}_t)$ is indeed higher than the volatility $\sigma(\widehat{p}_t^*)$ that would have been expected conditional on dividends across all 50 countries, with the only exception of Spain (where the volatility ratio is equal to 0.5). These volatility ratios are consistent with the results in Shiller (1981), suggesting that the earlier findings of excess volatility in the US can be extended in an international context.

Interestingly, the magnitude of excess volatility in many countries appears to be considerably higher compared to the commonly quoted ratio of 5:1 in the US market that was reported in Shiller (1981). For instance, the volatility ratio θ ranges from a minimum of 0.5 (Spain) to a maximum of 22.7 (Malaysia), with a mean (median) value of 7.5 (6.0). Overall, in addition to the ratio almost universally exceeding one (i.e. $\sigma(\widehat{p}_t)$ exceeding $\sigma(\widehat{p}_t^*)$), the volatility ratio exceeds the value of 5 (reported in the US by Shiller, 1981) for 27 out of 50 sample countries.

[Figure 4 about here.]

5.3 Second Generation Bounds Tests

The dividends' order of integration could potentially explain the empirical observation about the time-series of ex-post rational prices appearing much smoother than the time-series of realized prices. In this sense, the apparent smoothness of the ex-post rational price time-series does not necessarily imply that its variance is lower than that of realized prices when dividends are persistent. Engel (2005), in particular, states that "*Given the near-random-walk behavior of stock prices, the volatility of the stock price is captured by* $var(p_t - p_{t-1})$", in which case the relatively high volatility of the realized price would be consistent with the reversed variance inequality in (15).

Since dividends seem to follow a unit root process in our sample countries, we proceed by evaluating the Engel (2005) bounds test on the variance of first differences in prices,

[12] Tables D1 - D5 in the Internet Appendix present more detailed results of the first generation Shiller (1981) bounds tests. For instance, we find that the magnitude of excess volatility seems to be inversely related to the length of the available time-series. Moreover, the volatility ratios associated with inequalities (13) and (14) about price change variance also exceed the value of unity in all sample countries, consistent with the results in Shiller (1981) for the US. For robustness, we also perform those tests when expected dividend growth has been obtained via regressions of dividend growth against the dividend-price ratio, using annually-aggregated dividends (as in Cochrane, 2008) or monthly dividends in MIDAS regressions (as in Asimakopoulos et al., 2017).

given in inequality (15). Table 6 reports a set of summary statistics of the volatility ratio $\frac{\Delta\sigma(\widehat{p}_t)}{\Delta\sigma(\widehat{p}_t^*)}$ across the 50 sample countries. Panel A reports the results when expected dividend growth is obtained via the Cochrane-type regressions at an annual frequency in (2), while Panel B reports the results when expected dividend growth is obtained via the MIDAS regressions in (3).[13]

The results fail to reject the hypothesis of market efficiency when we account for dividends' non-stationarity via the second generation Engel (2005) variance bounds test. Irrespective of the specific measure of expected dividend growth used, the volatility ratio exceeds unity across all sample countries (with the single exception of the Czech Republic, where the volatility ratio takes the fairly borderline value of 0.97 when expected dividend growth is computed via MIDAS regressions). In other words, the volatility of first differences in realized prices consistently exceeds that of first differences in ex-post rational prices, consistent with what would have been expected if markets are efficient.

Importantly, the resulting volatility ratios $\frac{\sigma(\Delta\widehat{p}_t)}{\sigma(\Delta\widehat{p}_t^*)}$ are markedly different depending on the way in which expected dividend growth has been obtained. When we use regressions at an annual frequency (as in Cochrane, 2008), the mean volatility ratio is equal to 14.1, and it ranges from a minimum of 4.7 (Romania) to a maximum of 60.2 (Argentina). In contrast, estimating expected dividend growth via MIDAS regressions (as in Asimakopoulos et al., 2017) results in substantially lower volatility ratios across all countries. In the latter case, the mean volatility ratio is equal to 2.8 and it ranges from a minimum of 1.0 (Czech Republic) to a maximum of 14.5 (Taiwan). In general, employing the MIDAS framework to obtain a measure of expected dividend growth appears to result in volatility ratios that are lower by a factor of around 5 compared to the ones obtained under Cochrane-type regressions at an annual frequency.

[Table 6 about here.]

Overall, both approaches for obtaining expected dividend growth result in volatility ratios that do not violate the Engel (2005) inequality and, thus, do not reject the hypothesis of efficient markets. However, the MIDAS framework appears to result in $\sigma(\Delta\widehat{p}_t)$ being substantially closer to $\sigma(\Delta\widehat{p}_t^*)$, compared to the alternative approach of Cochrane-type regressions at an annual frequency. Hence, we provide evidence that accounting for dividend seasonality results in the variance of realized price changes being much closer to the value that would have been expected based on the subsequent dividends. This large impact of the mixed-frequency estimation on the volatility ratios is also evident from Figure 5, which illustrates how the ratios produced by the MIDAS approach are much closer to the theoretical threshold of unity compared to those produced by Cochrane-type regressions at an annual frequency.

[13] The individual volatility ratios on a country-by-country basis are reported in Table E1 in the Internet Appendix, to conserve space.

[Figure 5 about here.]

Another interesting finding refers to an observed negative relationship between the volatility ratio and the length of available data for each country when expected dividend growth is obtained via Cochrane-type regressions. As can be seen from Panel A of Table 6, countries with longer time-series of available data tend to to have substantially lower volatility ratios compared to countries with relatively shorter time-series of available data, although this relationship is not strictly monotonic. For example, Canada (with the longest time-series of available data at 187 years) has a volatility ratio of 8.5 while, at the other end of the spectrum, Bulgaria (with the shortest time-series of available data at 17 years) has a markedly higher volatility ratio of 32.9. More generally, the mean volatility ratio is equal to 12.4 in the subsample of countries with more than 60 years of available data, while the respective mean is equal to 15.3 for countries with fewer than 30 years of available data.[14]

In contrast, this negative relationship between sample length and the volatility ratio largely disappears when we use the MIDAS framework to obtain a measure of expected dividend growth (Panel B of Table 6). More specifically, in this case the mean volatility ratio is equal to 2.3 for countries with fewer than 30 years of available data and 2.2 for countries with more than 60 years of data.

6 Conclusion

The simple notion that rational investors determine current stock prices as the sum of discounted expected dividends has attracted a lot of attention in the empirical literature. We find that the dividend-price ratio has markedly higher predictive power over subsequent dividend growth when used at a monthly frequency in MIDAS regressions compared to the standard practice of using annually-aggregated values. Importantly, the variance of realized price changes is consistently closer to that of expected price changes when using the MIDAS approach, indicating that dividend seasonality can potentially explain a substantial part of the previously reported excess volatility puzzle.

We show that the results of variance bounds tests of market efficiency ultimately depend on the assumptions one is willing to make about the underlying dividend process and on whether dividend data are aggregated when used to predict dividend growth. Thus, we find that assuming stationary dividends results in the detection of excess volatility across all our sample countries. However, given that in reality dividends deviate substantially

[14] The negative relationship between the volatility ratio and the sample length is even more pronounced when we estimate the ratio based on the Shiller (1981) inequality in (12). In this case, countries with fewer than 30 years of available data have a mean volatility ratio of 12.5, while countries with more than 60 years of data have a mean ratio of 2.4. Tables D1 - D5 in the Internet Appendix report more detailed results about the relationship between excess volatility and sample length for the case of first generation bounds tests.

from stationarity, these results seem to say more about the assumption's validity than about market efficiency. Second generation bounds tests, which allow for non-stationary dividends, fail to reject the hypothesis of efficient markets in any of the sample countries, casting doubt on the existence of the excess volatility puzzle.

References

Almon, S. (1965). The distributed lag between capital appropriations and expenditures. *Econometrica: Journal of the Econometric Society*, pages 178–196.

Andreou, E., Ghysels, E., and Kourtellos, A. (2013). Should macroeconomic forecasters use daily financial data and how? *Journal of Business & Economic Statistics*, 31(2):240–251.

Ang, A. and Bekaert, G. (2007). Stock return predictability: Is it there? *Review of Financial Studies*, 20:651–707.

Asgharian, H., Hou, A. J., and Javed, F. (2013). The importance of the macroeconomic variables in forecasting stock return variance: A GARCH-MIDAS approach. *Journal of Forecasting*, 32(7):600–612.

Asimakopoulos, P., Asimakopoulos, S., Kourogenis, N., and Tsiritakis, E. (2017). Time-disaggregated dividend-price ratio and dividend growth predictability in large equity markets. *Journal of Financial and Quantitative Analysis*, 52:2305–2326.

Bai, J., Ghysels, E., and Wright, J. H. (2013). State space models and MIDAS regressions. *Econometric Reviews*, 32(7):779–813.

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

Clements, M. P. and Galvão, A. B. (2008). Macroeconomic forecasting with mixed-frequency data: Forecasting output growth in the United States. *Journal of Business & Economic Statistics*, 26(4):546–554.

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

Engel, C. (2005). Some new variance bounds for asset prices. *Journal of Money, Credit and Banking*, 37:949–955.

Forsberg, L. and Ghysels, E. (2007). Why do absolute returns predict volatility so well? *Journal of Financial Econometrics*, 5(1):31–67.

Gagliardini, P., Ghysels, E., and Rubin, M. (2017). Indirect inference estimation of mixed frequency stochastic volatility state space models using MIDAS regressions and ARCH models. *Journal of Financial Econometrics*, 15(4):509–560.

Ghysels, E., Santa-Clara, P., and Valkanov, R. (2005). There is a risk-return trade-off after all. *Journal of Financial Economics*, 76(3):509–548.

Ghysels, E., Santa-Clara, P., and Valkanov, R. (2006). Predicting volatility: getting the most out of return data sampled at different frequencies. *Journal of Econometrics*, 131(1-2):59–95.

Ghysels, E., Sinko, A., and Valkanov, R. (2007). MIDAS regressions: Further results and new directions. *Econometric Reviews*, 26(1):53–90.

Golez, B. and Koudijs, P. (2018). Four centuries of return predictability. *Journal of Financial Economics*, 127:248–263.

Lansing, K. (2016). On variance bounds for asset price changes. *Journal of Financial Markets*, 28:132–148.

Lansing, K. and LeRoy, S. (2014). Risk aversion, investor information and stock market volatility. *European Economic Review*, 70:88–107.

LeRoy, S. and Porter, R. (1981). The present-value relation: Tests based on implied variance bounds. *Econometrica*, 49:555–574.

Marsh, T. and Merton, R. (1986). Dividend variability and variance bounds tests for the rationality of stock market prices. *American Economic Review*, 76:483–498.

Merton, R. (1986). On the current state of the stock market rationality hypothesis.

Rangvid, J., Schmeling, M., and Schrimpf, A. (2014). Dividend predictability around the world. *Journal of Financial and Quantitative Analysis*, 49:1255–1277.

Shiller, R. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71:421–436.

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

West, K. (1988). Dividend innovations and stock price volatility. *Econometrica*, 56:37–61.

Table 1: Number of countries rejecting the stationarity and unit root tests.

	$\widehat{d}_t$			$\widehat{p}_t$		
	ADF test $H_0 : I(1)$	PP test $H_0 : I(1)$	KPSS $H_0 : I(0)$	ADF test $H_0 : I(1)$	PP test $H_0 : I(1)$	KPSS $H_0 : I(0)$
n rejecting H_0 at $\alpha = 1\%$	1	1	24	4	6	19
n rejecting H_0 at $\alpha = 5\%$	4	5	35	6	9	32
n rejecting H_0 at $\alpha = 10\%$	8	9	50	8	10	50
n	50		50	50	50	50

Notes: This Table presents the results from a set of unit root and stationarity tests on the time-series of dividends $\widehat{d}_t$ and stock prices $\widehat{p}_t$. The Table reports the number of sample countries that reject the null hypothesis in each test, namely in the Augmented Dickey Fuller (ADF) test, the Phillips-Perron (PP) test, and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test. The null hypothesis in the ADF test and the PP test is that the respective time-series contains a unit root, while the null hypothesis in the KPSS test is that the time-series is stationary. Each test is performed separately per country, across a sample of 50 countries.

Table 2: Pool unit root tests on prices and dividends

Panel A: Prices and dividends					
	Null	Results			
Test	Hypothesis	p	d	$p_t - p_{t-1}$	$d_t - d_{t-1}$
lm-Pesaran-Shin W	Unit root	-1.78	-0.19	-36.90	-33.46
		(0.04)	(0.43)	(0.00)	(0.00)
Augmented Dickey-Fuller Fisher χ^2	Unit root	4.61	3.19	114.32	101.71
		(0.00)	(0.00)	(0.00)	(0.00)
Phillips-Perron Fisher χ^2	Unit root	4.76	2.68	154.22	128.85
		(0.00)	(0.00)	(0.00)	(0.00)
Panel B: De-trended prices and dividends					
	Null	Results			
Test	Hypothesis	p	d	$p_t - p_{t-1}$	$d_t - d_{t-1}$
lm-Pesaran-Shin	Unit root	-9.90	-7.10	-35.84	-33.10
		(0.00)	(0.00)	(0.00)	(0.00)
Augmented Dickey-Fuller Fisher χ^2	Unit root	17.00	10.15	111.12	99.57
		(0.00)	(0.00)	(0.00)	(0.00)
Phillips-Perron Fisher χ^2	Unit root	18.75	9.14	151.10	137.67
		(0.00)	(0.00)	(0.00)	(0.00)

Notes: This Table presents the results of a set of pool unit root tests for the time-series of prices and dividends. The pool unit root tests include the Im-Persaran-Shin, the Augmented Dickey-Fuller, and Phillips-Perron test. Panel A refers to the original price and dividend series, while Panel B refers to the respective de-trended time-series. The results are from a panel dataset for the 50 countries included in the study. The Table reports the test statistics and the respective p-values (in brackets).

Table 3: Performance of MIDAS weighting schemes

	Step function	Exponential Almon	Beta	Beta constrained	Almon P
mean RMSE	0.079	0.105	0.104	0.103	0.070
median RMSE	0.062	0.067	0.074	0.062	0.065
min RMSE	0.000	0.002	0.003	0.002	0.000
max RMSE	0.542	0.714	0.488	0.503	0.255
n selected	8	8	8	6	20
n	50	50	50	50	50

Notes: The Table reports the Root Mean Squared Error of alternative MIDAS weighting schemes in regressions of (annual) changes in dividend growth against the lagged (monthly) dividend-price ratio. The weighting schemes include a polynomial with a step-function, an exponential Almon lag polynomial, a normalized beta polynomial, a normalized beta polynomial where the last lag is constrained to zero, and an Almon lag polynomial of order P. The MIDAS regressions are estimated separately for each country, across a sample of 50 countries. For each country, we estimate each candidate scheme using an estimation period that starts at the beginning of the full sample period and ends 2 years before the end of the sample period. Then, we use the in-sample model parameters to produce out-of-sample forecasts for the last 2 years of the full sample period. The Table reports statistics (mean, median, minimum and maximum) of the RMSE of the out-of-sample forecasts associated with each scheme, as well as the number of countries for which a particular scheme has been selected based on having the lowest RMSE.

Table 4: Predictive regressions for dividend growth

Country	n	Shiller-type regressions slope	Shiller-type regressions R^2	Cochrane-type regressions slope	Cochrane-type regressions R^2	MIDAS regressions slope	MIDAS regressions R^2
Bulgaria	17	0.51***	0.50	0.42***	0.14	0.09	0.90
Romania	19	0.64***	0.53	0.21*	0.09	0.07	0.85
Russia	20	0.58***	0.56	0.01	0.00	0.06	0.78
Tunisia	21	0.57***	0.55	0.13	0.04	0.05	0.74
Brazil	22	0.64***	0.56	0.10	0.01	0.05	0.53
Czech	23	0.44***	0.51	0.31**	0.15	0.05	0.96
Hungary	23	0.54***	0.52	-0.09*	0.03	0.02	0.55
Poland	23	0.48***	0.50	0.06	0.01	0.06	0.88
Israel	24	0.4***	0.57	0.21	0.07	0.08	0.70
Egypt	25	0.36***	0.58	0.03	0.02	0.04	0.86
China	27	0.41***	0.55	-0.02	0.00	0.07	0.35
Indonesia	27	0.43***	0.52	-0.25	0.05	0.08	0.75
Ireland	27	0.45***	0.52	0.01	0.00	0.05	0.67
Portugal	31	0.41***	0.48	-0.05	0.01	0.04**	0.44
Colombia	32	0.40***	0.58	0.21*	0.04	0.07	0.64
Nigeria	32	0.49***	0.51	0.40**	0.20	0.05**	0.49
Taiwan	32	0.40***	0.52	0.13	0.06	0.07	0.72
Turkey	32	0.51***	0.52	0.15**	0.10	0.03***	0.44
Kenya	34	0.38***	0.45	0.16**	0.06	0.07*	0.58
Morocco	34	0.41***	0.59	-0.04	0.03	0.01	0.31
Philippines	36	0.34***	0.53	0.02	0.00	0.06	0.45
Jordan	39	0.31***	0.52	0.16*	0.06	0.06	0.68
Greece	41	0.26***	0.49	0.13**	0.08	0.05	0.62
Thailand	42	0.24***	0.53	0.08	0.02	0.06	0.56
Chile	44	0.31***	0.65	0.29**	0.06	0.10***	0.36
Malaysia	44	0.24***	0.56	0.03	0.00	0.03	0.32
Singapore	45	0.25***	0.57	0.05**	0.05	0.02	0.23
Norway	48	0.17***	0.58	0.08*	0.04	0.04	0.60
Hong Kong	53	0.26***	0.60	0.15**	0.06	0.05	0.49
South Africa	54	0.27***	0.55	0.06	0.01	0.03	0.37
Korea	55	0.19***	0.58	0.05*	0.04	0.06***	0.54
Finland	56	0.22***	0.61	-0.09**	0.07	0.02**	0.66
Argentina	71	0.26***	0.50	0.01	0.00	0.06	0.45
United Kingdom	84	0.14***	0.55	0.06**	0.04	0.01*	0.16
New Zealand	91	0.11***	0.49	0.08	0.01	0.06***	0.64
Austria	93	0.10***	0.53	-0.10**	0.04	0.04	0.16
Italy	93	0.12***	0.45	0.15	0.04	0.08***	0.92
India	97	0.14***	0.52	-0.02	0.00	0.06***	0.65
Switzerland	99	0.08***	0.61	-0.07	0.02	0.08	0.58
Sweden	116	0.05***	0.64	-0.10*	0.04	0.05*	0.62
Japan	117	0.03***	0.33	-0.05**	0.05	0.02*	0.09
Spain	118	0.08***	0.45	-0.02	0.00	0.02*	0.16
Netherlands	126	0.07***	0.50	0.01	0.00	0.06**	0.58
Denmark	144	0.05***	0.45	-0.03	0.01	0.05***	0.52
Belgium	147	0.09***	0.43	0.09*	0.02	0.07**	0.74
United States	147	0.06***	0.63	-0.03*	0.02	0.03**	0.25
Germany	148	0.08***	0.20	-0.08	0.02	0.07	0.31
Australia	156	0.08***	0.58	0.00	0.00	0.05**	0.39
France	178	0.06***	0.48	-0.02	0.00	0.04**	0.34
Canada	187	0.07***	0.61	0.12**	0.04	0.07*	0.73

Notes: The Table reports the estimated intercept and slope coefficients, and the R-square from predictive least squares regressions for dividend growth. Statistical significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively. The first panel reports the results from regressing stock prices against time, as in Shiller (1981). The second panel reports the results from regressing dividend growth against the lagged dividend-price ratio at an annual frequency (as in Cochrane, 2008). The third panel reports the results from mixed-frequency data sampling (MIDAS) regressions of dividend growth against the lagged dividend-price ratio. Countries are sorted in ascending order based on the number of years in second column with available data (n).

Table 5: In-sample R^2 of predictive regressions for dividend growth

	Shiller-type regressions	Cochrane-type regressions	MIDAS regressions
mean	0.53	0.04	0.55
median	0.53	0.04	0.57
min	0.20	0.00	0.09
max	0.65	0.20	0.96

Notes: The Table reports descriptive statistics for the R-square obtained from predictive least squares regressions for dividend growth. The second column refers to results from regressing stock prices against time, as in Shiller (1981). The third column refers to results from regressing dividend growth against the lagged dividend-price ratio at an annual frequency (as in Cochrane, 2008). The fourth column refers to results from mixed-frequency data sampling (MIDAS) regressions of dividend growth against the lagged dividend-price ratio. Each regression is estimated separately for each country, across a sample of 50 countries. The Table reports the mean, median, minimum and maximum estimated R-square across the 50 sample countries.

Table 6: Variance bounds tests - Engel volatility ratio

Panel A: Expected growth via Cochrane-type regressions				
	Full sample	Subsamples formed on time-series length		
		$L \leq 30$	$30 < L \leq 60$	$60 < L$
mean	14.1	15.3	15.0	12.4
median	10.5	10.5	11.6	9.2
min	4.7	4.7	5.5	4.7
max	60.2	32.9	58.1	60.2
$n > 1$	50	13	19	18
n	50	13	19	18
Panel B: Expected growth via MIDAS regressions				
	Full sample	Subsamples formed on time-series length		
		$L \leq 30$	$30 < L \leq 60$	$60 < L$
mean	2.8	2.3	3.6	2.2
median	2.2	2.4	2.8	1.7
min	1.0	1.0	1.2	1.0
max	14.5	4.4	14.5	9.0
$n > 1$	49	12	19	18
n	50	13	19	18

Notes: This Table presents the results of computing the Engel (2005) volatility ratio $\frac{\sigma(\Delta \widehat{p}_t)}{\sigma(\Delta \widehat{p}_t^*)}$. Panel A reports the results when expected dividend growth is obtained via Cochrane-type regressions of changes in dividend growth against the lagged dividend-price ratio at an annual frequency, while Panel B reports the results when expected dividend growth is obtained via MIDAS regressions of changes in (annual) dividend growth against the lagged (monthly) dividend-price ratio. The Table reports the mean, median, minimum, and maximum values of each ratio across the 50 sample countries, the number of cases where the ratio exceeds the value of 1, and the number of countries n. The first column of each panel reports the results across the full sample of 50 countries, while the last three columns report the results from subsamples that have been formed based on the countries' number of years of available data L.

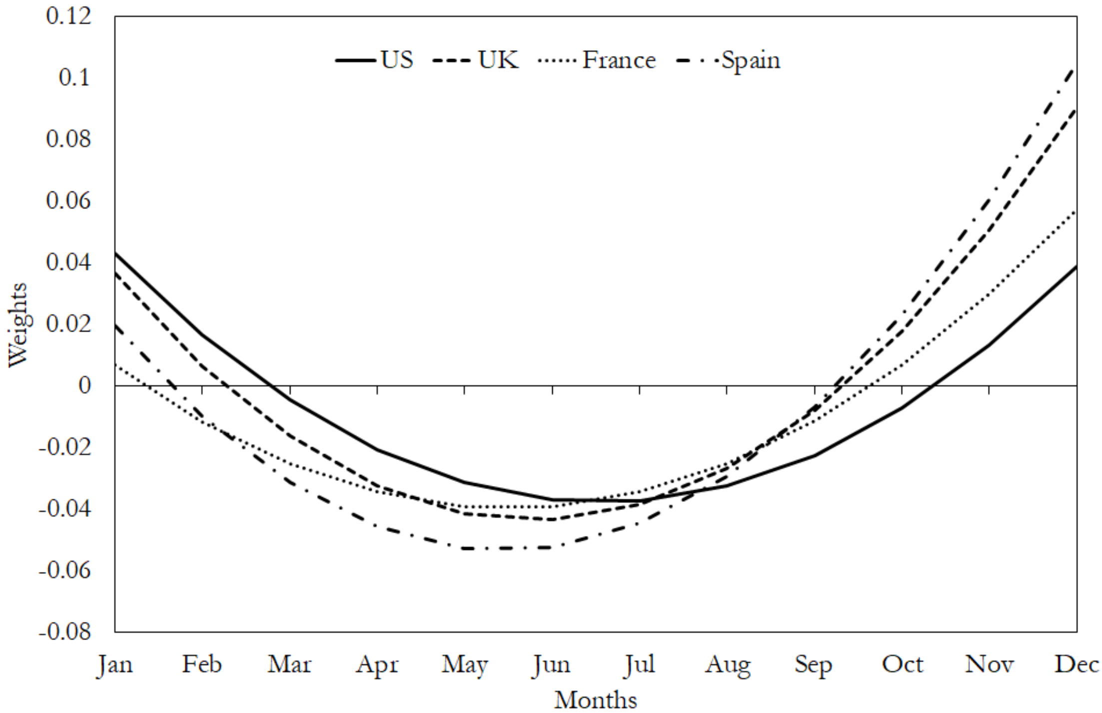

Figure 1: MIDAS weighting scheme (Almon P)

Notes: This Figure plots the MIDAS weighting scheme from regressions of (annual) changes in dividend growth against the lagged (monthly) dividend-price ratio. The Figure plots the associated weights for a subsample of four countries, namely the US, UK, France, and Spain. The MIDAS estimation is based on the Almon lag polynomial of order P.

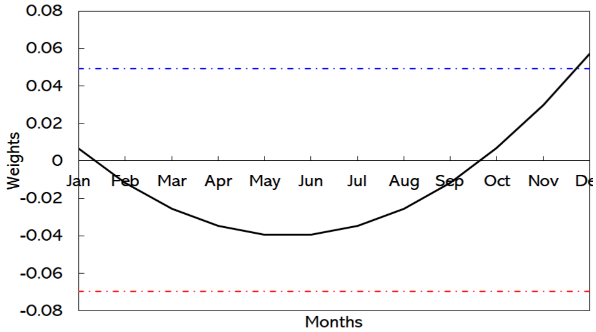

(a) France

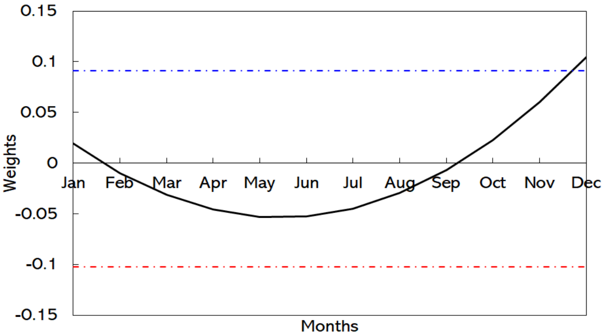

(b) Spain

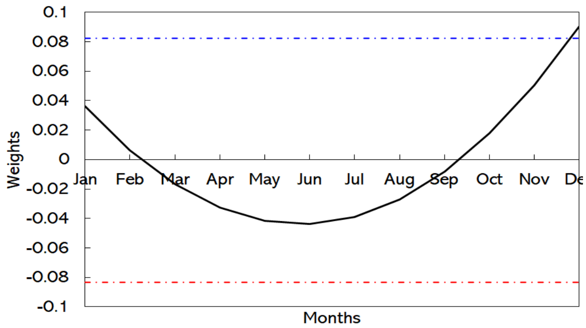

(c) UK

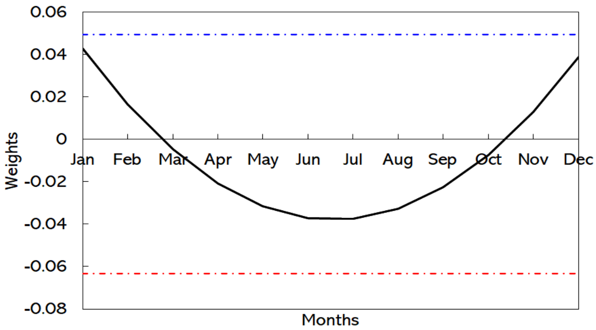

(d) US

Figure 2: MIDAS weighting scheme (Almon P)

Notes: This Figure plots the MIDAS weighting scheme from regressions of (annual) changes in dividend growth against the lagged (monthly) dividend-price ratio. The Figure plots the associated weights for a subsample of four countries, namely France, Spain, UK, and US. The MIDAS estimation is based on the Almon lag polynomial of order P. The lower and upper 95% confidence bounds of the weights are respectively illustrated as the red and blue dashed lines.

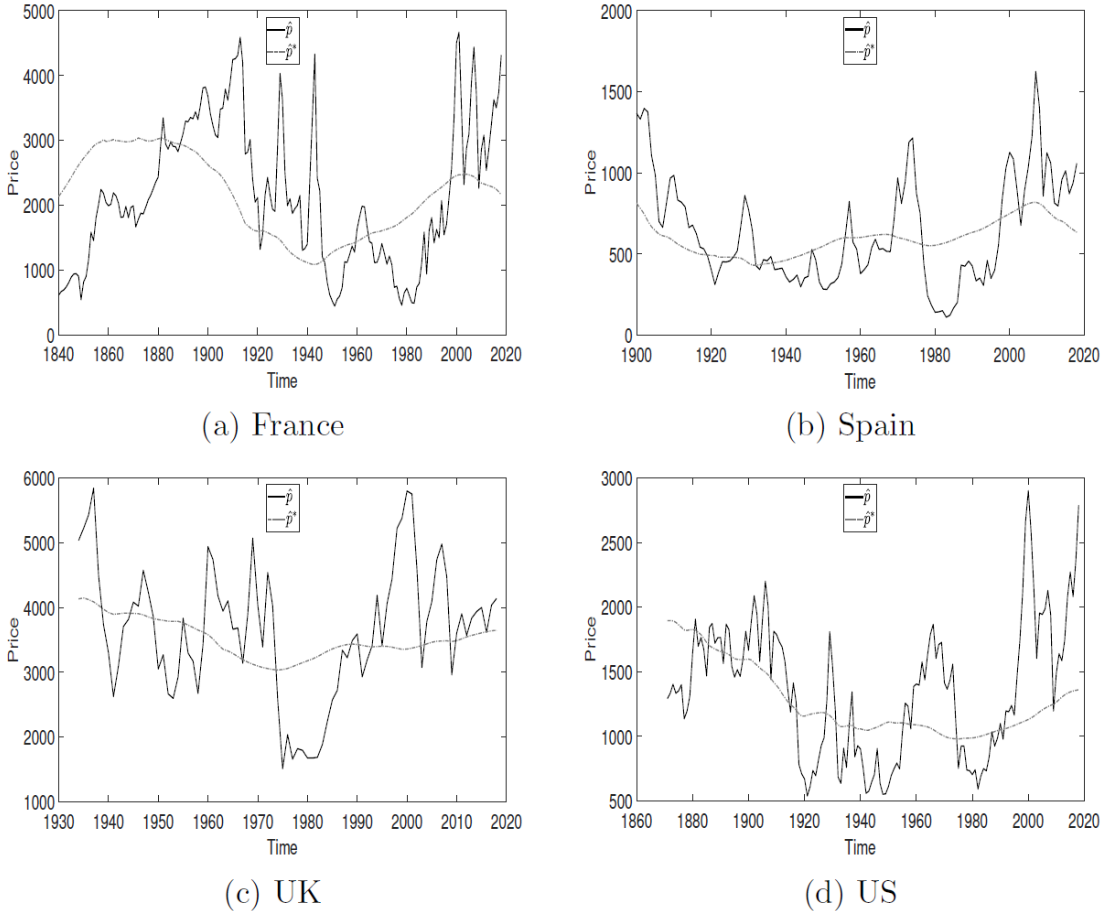

Figure 3: Time-evolution of realized price and ex-post rational price

Notes: This Figure plots the time-series of de-trended realized prices $\widehat{p}_t$ and de-trended ex-post rational prices $\widehat{p}_t^*$ for a subsample of four countries, namely France, Spain, UK and US.

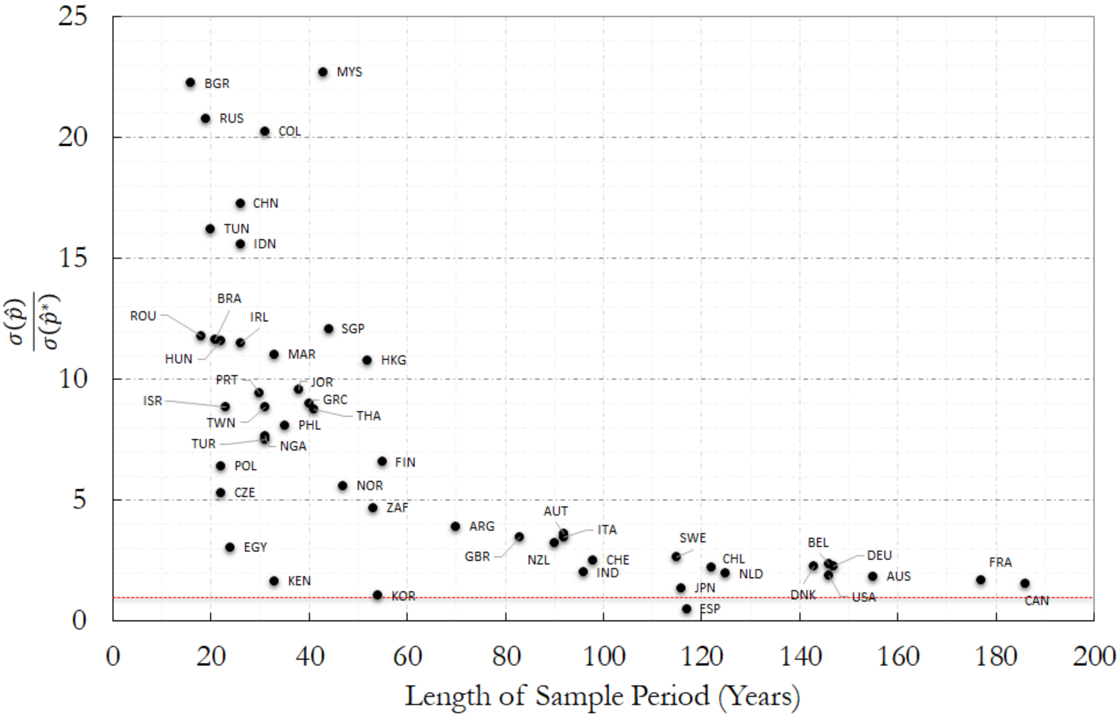

Figure 4: Shiller volatility ratios and sample period length

Notes: This Figure plots the Shiller (1981) ratio $\frac{\sigma(\widehat{p}_t)}{\sigma(\widehat{p}_t^*)}$ of the volatility of the realized price over the volatility of the ex-post rational price, across a sample of 50 countries. The volatility ratios have been ordered (on the horizontal axis) in ascending order based on the number of years of available data for each country. The horizontal line indicates the threshold value of 1. Ratios exceeding the value of 1 violate the Shiller (1981) variance bound.

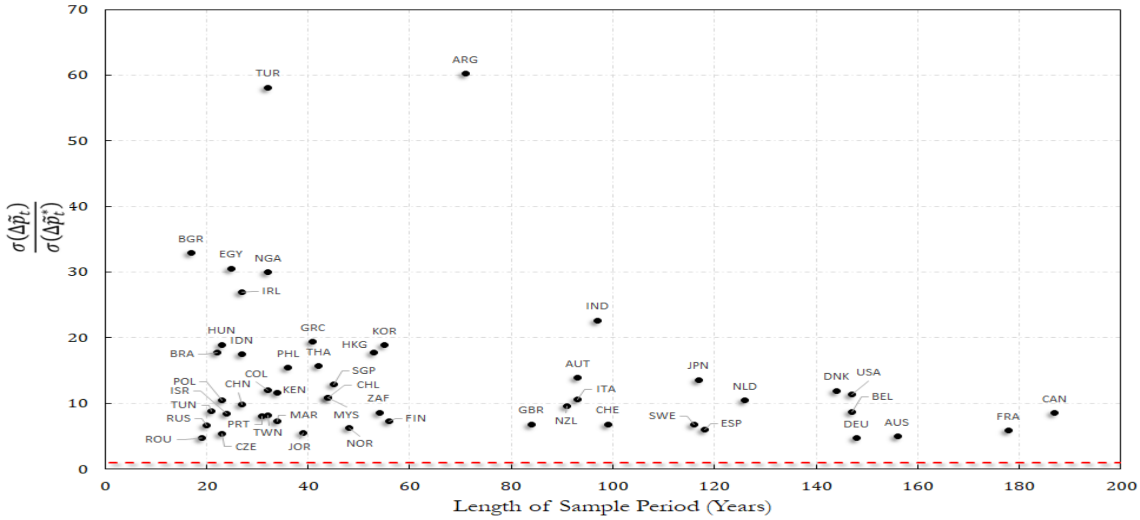

(a) Cochrane-type regressions

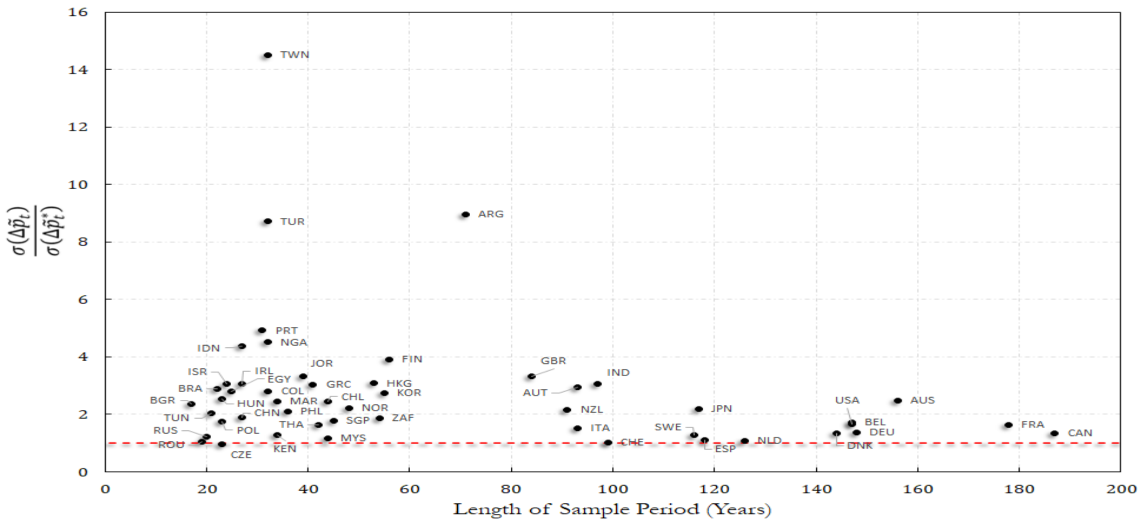

(b) MIDAS regressions

Figure 5: Engel volatility ratios and sample period length

Notes: This Figure plots the Engel (2005) ratio $\frac{\sigma(\Delta\widehat{p}_t)}{\sigma(\Delta\widehat{p}_t^*)}$ of the volatility of realized price changes over the volatility of ex-post rational price changes, across a sample of 50 countries. The volatility ratios have been ordered (on the horizontal axis) in ascending order based on the number of years of available data for each country. The horizontal line indicates the threshold value of 1. Ratios falling below the value of 1 violate the Engel (2005) variance bound. Subfigure (a) plots the volatility ratios when expected dividend growth is obtained via Cochrane-type regressions of changes in dividend growth against the lagged dividend-price ratio at an annual frequency, while subfigure (b) plots the respective ratios when expected dividend growth is obtained via MIDAS regressions of changes in (annual) dividend growth against the lagged (monthly) dividend-price ratio.