Received: 1 May 2023 | Accepted: 20 April 2024

DOI: 10.1111/jfir.12403

The Journal of Financial Research

ORIGINAL ARTICLE

MIDAS and dividend growth predictability: Revisiting the excess volatility puzzle

Enoch Quaye[1] | Radu Tunaru[2] | Nikolaos Voukelatos[3]

[1]School of Accounting and Finance, University of Bristol, Bristol, UK

[2]University of Sussex Business School, University of Sussex, Brighton, UK

[3]Kent Business School, University of Kent, Canterbury, UK

Correspondence
Nikolaos Voukelatos, Kent Business School, University of Kent, Canterbury CT2 7FS, UK.
Email: N.Voukelatos@kent.ac.uk

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 nonstationary dividends consistently fail to reject the market efficiency hypothesis across all countries. Our findings suggest that the common rejection of market efficiency in the literature is most likely driven by the annual aggregation of dividend data as well as by the assumption of stationary dividends.

JEL CLASSIFICATION
G12, G14, G17

1 | INTRODUCTION

There are few topics in finance and economics that have attracted as much attention as the question of what determines stock prices and, in particular, whether stock prices reflect the fundamental value of the underlying firms. The debate around this question has been running the full spectrum from Keynes's views about stock markets operating as casinos for the lucky to the efficient market hypothesis of Fama and Samuelson. One approach that has been adopted to evaluate the rationality of stock markets refers to the variance bounds tests 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 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 is 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 & Bekaert, 2007; Cochrane, 2008; van Binsbergen & 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 article, we adopt a new approach for evaluating market efficiency by accounting for intrayear 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 April–June. Although 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 article also contributes to the literature on dividend growth predictability and excess volatility by shifting the focus from the United States to a large sample of international stock markets. Studies have tended to focus almost exclusively on the US stock market. Although this emphasis on the United States 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, revisiting 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 and 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 to use MIDAS regressions in this context.

Under this new modeling 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 for 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 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 variance bounds tests. We show that the results of variance bounds tests are highly sensitive to the approach used for obtaining expected dividend growth, highlighting the importance of accounting for intrayear dividend seasonality.

We also report evidence of varying performance for different weighting schemes in the MIDAS regressions. In particular, we consider a set of different MIDAS weighting schemes and base our selection of the optimal scheme for each country on an out-of-sample forecasting exercise. We find that in most 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 toward the end of the year take substantially higher weights in the MIDAS estimations compared to dividends that are paid 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 United States 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 most cases, the ratio of realized price volatility over the volatility of ex post rational prices exceeds the value of 5:1 reported in Shiller (1981) for the US market, suggesting that deviations from the market rationality hypothesis could 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 is in fact reversed.[1]

We find that the assumption of stationary dividends is consistently violated in our international sample, and we 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 market efficiency hypothesis. 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 studies is most likely driven by dividend nonstationarity and the discarding of information from the annual aggregation of dividends, as opposed to reflecting genuine market inefficiencies.

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 (2148 months). However, given that coverage in GFD begins at different times for different countries, the number of available observations varies across sample countries. For example, data for the French equity market is available from 1840 whereas, at the other end of the spectrum, coverage for Bulgaria begins only in 2001. Finally, the fact that we focus on aggregate dividends at the stock index level means that there are no gaps in the time series of dividends, with observations of aggregate stock index dividends available in every month for every country after the beginning of its respective sample period.

The MIDAS estimation involves regressing the dividend growth (at an annual frequency) against the lagged first difference of the dividend–price ratio (at a monthly frequency). The monthly quoted variables $P_{t,k}^m$ and $D_{t,k}^m$ denote the price levels and dividend levels, respectively, observed in year t and month k. The variables $p_{t,k}^m$ and $d_{t,k}^m$ 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 (except for Egypt). Despite its almost universally positive sign, mean dividend growth varies substantially across countries, ranging from 0.13 (Japan) to 0.465 (Argentina). Interestingly, countries with shorter available data series generally

[1]More recently, Lansing and LeRoy (2014) demonstrate that the relation between log-return variance and investor information can be nonmonotonic, depending on the 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 Table 4 in Section 5 for the full list of countries in our article.

[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.

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

	$\hat{d}_t$			$\hat{p}_t$		
	ADF	PP	KPSS	ADF	PP	KPSS
	H0: $I(1)$	H0: $I(1)$	H0: $I(0)$	H0: $I(1)$	H0: $I(1)$	H0: $I(0)$
n rejecting H0 at α = 1%	1	1	24	4	6	19
n rejecting H0 at α = 5%	4	5	35	6	9	32
n rejecting H0 at α = 10%	8	9	50	8	10	50
n	50	50	50	50	50	50

Note: This table presents the results from a set of unit root and stationarity tests on the time series of dividends $\hat{d}_t$ and stock prices $\hat{p}_t$. The table reports the number of sample countries that reject the null hypothesis in each of the following tests: augmented Dickey Fuller (ADF) test, Phillips–Perron (PP) test, and 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, whereas 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.

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 has only 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 (except for Portugal).[4]

We perform 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, with the null of stationarity being 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 similar, with the null of a unit root being rejected at the 5% level in only 5 countries (Taiwan, New Zealand, Italy, Australia, and Canada).

We also perform a set of pool unit root tests on the original and detrended time series of prices and dividends. More specifically, Table 2 reports the results of the Im–Persaran–Shin, ADF, and PP unit root tests using pooled data across all 50 sample countries. Overall, the results reject the hypothesis of a unit root in the pooled data set.

3 | MIDAS AND EXPECTED DIVIDEND GROWTH

3.1 | Expected dividend growth

Because 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.

[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 comovement between the two types of time series. Descriptive statistics on the time series of index returns, dividend growth, and dividend yield in each country are reported in Table A1 in the Online Appendix.

[5]The results of the stationarity and unit root tests are reported in Table B1 in the Online Appendix.

TABLE 2 Pool unit root tests on prices and dividends.

Test	Null hypothesis	Results p	d	$p_t - p_{t-1}$	$d_t - d_{t-1}$
Panel A: Prices and dividends					
Im–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	114.22	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: Detrended prices and dividends					
Im–Pesaran–Shin W	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)

Note: 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, augmented Dickey–Fuller, and Phillips–Perron tests. Panel A refers to the original price and dividend series, and Panel B refers to the respective detrended time series. The results are from a panel data set for the 50 countries included in the study. The table reports the test statistics and their respective p-values (in parentheses).

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:

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

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. 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 Asimakopoulos et al., 2017; Cochrane, 2008; Golez & 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:

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

where subscript t denotes time at an annual frequency. Following Cochrane (2008), the standard errors are generalized method of moments (GMM) corrected for heteroskedasticity. Similar to Ang and Bekaert (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 detrending 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) leads to some loss of within-year information. We address this concern by employing a MIDAS framework, where lower frequency (annual) dividend growth Δg_{t+1} is regressed against the higher frequency (monthly) first difference of the dividend–price ratio Δy_t^m.

3.2 | MIDAS framework

The MIDAS approach is introduced by Ghysels et al. (2005, 2006) as a framework to estimate regression specifications where the dependent and independent variables are sampled at different frequencies. Subsequent empirical studies use 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, whereas Clements and Galvão (2008), Bai et al. (2013), and Gagliardini et al. (2017) use monthly macroeconomic and financial indicators to forecast quarterly gross domestic product (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 first difference of the dividend–price ratio over subsequent annual dividend growth.

We follow Asimakopoulos et al. (2017) and apply the MIDAS framework to estimate the effect of the higher frequency data of the first difference of the log dividend yield Δy_t^m (monthly, $m = 12$) on the lower frequency data of the log dividend growth rate Δg_{t+1} (annual). The applied regression model is

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

for $t = 1, \ldots, T$, where $L^{1/m}$ denotes a lag operator. 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, and $L^{k/m}$ is the lag operator of Δy_t^m for k/m periods. The maximum length of the polynomial function is $K - 1$. The overall impact of the lagged Δy_t^m on Δg_{t+1} is given by the coefficient $\hat{\beta}_{1,g}$, which can be obtained by normalizing the weights $\omega_k(\boldsymbol{\theta})$ so they sum to one.[6] The general form Equation (3) can be rewritten as

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

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

3.3 | MIDAS weighting schemes

The shape of the weighting scheme ω depends on the chosen specification for the polynomial function. On this issue, Ghysels et al. (2007) discuss alternative weighting schemes. In this article, we consider four alternative

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

specifications for the polynomial function, each of which results in a different weighting scheme and, by extension, a different effect of Δy_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 & Ghysels, 2007). Under this alternative, a regressor X_t can be expressed as the partial sum of the higher frequency x^m so that

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

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 $K_1 < \ldots < K_M$. The impact of x^m on the dependent variable can be measured by the sum of all coefficients in the OLS regression (i.e., $\sum_{i=1}^{M} \beta_i$) because it appears in all 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, similar 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

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

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 flexible, as it allows for weights that can take a variety of 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 are 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, we also consider a restricted version where the last lag is set to 0, with $\omega_k(\theta_1, \theta_2, 0)$.

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

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

The weights in Equation (9) are obtained via a nonlinear least squares estimation, where the optimal lag order is selected using the Akaike information criterion (AIC)/Bayesian information criteron (BIC) of the least squares estimation (for a more detailed discussion, see Pettenuzzo et al., 2016). 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.

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

Note: This table reports the root mean squared error (RMSE) of alternative mixed-frequency data sampling (MIDAS) weighting schemes in regressions of (annual) changes in dividend growth against the lagged (monthly) first differences of the 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 RMSE statistics (mean, median, minimum, and maximum) 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.

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 earlier. 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 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, whereas the exponential Almon polynomial produces the largest mean RMSE (0.105). However, the polynomial with a step function results in the lowest median RMSE across the 50 countries (0.062). More important, 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 of 50 countries. The step-function polynomial, unconstrained beta, and exponential Almon are jointly second (8 cases), and the constrained beta polynomial is selected the least (6 cases).[7]

All four weighting schemes that we consider for the MIDAS specification represent reasonable a priori candidates to capture the complex dynamics of the relation between the dividend–price ratio and subsequent dividend growth. Nevertheless, the relative ability of each scheme to capture that complex relation is, naturally, an

[7]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 for Bulgaria (with only 17 years of available data) and 0.03 for Canada (with 187 years of available data). Table C1 in the Online Appendix reports the RMSE per weighting scheme for each country.

empirical question. In this sense, our results suggest that no particular scheme can be considered as a universally optimal choice across all countries. Although the best performing Almon *P* polynomial is selected in most of the cases, the other four schemes are selected in 6–8 cases each, of the 50 sample countries.

Intuitively, our empirical results suggest that the relative simplicity of the Almon *P* polynomial, which can be estimated in a simple OLS setting with a linear transformation of high-frequency regressors, largely outweighs the benefits of more complex weighting schemes (see also Ghysels & Qian, 2019; Mogliani & Simoni, 2021; Pettenuzzo et al., 2016). For instance, the exponential Almon and the beta polynomials are generally considered flexible, to accommodate different weighting structures (see, e.g., Asimakopoulos et al., 2017; Ghysels et al., 2007; You & Liu, 2020). However, the positivity constraint imposed on the weights could prove too restrictive, as it rules out the case where the dividend–price ratio could positively affect dividend growth at certain lags and negatively at other lags (see also Breitung & Roling, 2015). In fact, Figure 1 shows that the Almon *P* polynomial regularly places negative weights at dividends paid in (or around) the summer months, a characteristic that the exponential Almon and the beta polynomials are not able to capture.

Another potential advantage of the Almon *P* polynomial over competing schemes is its ability to accommodate cases where the true weights decrease slowly across lags. In contrast, alternative schemes based on beta functions or exponential Almon polynomials are better suited to accommodate rapidly declining weights (Ghysels & Qian, 2019). In this sense, our finding that the Almon *P* polynomial largely outperforms competing schemes could

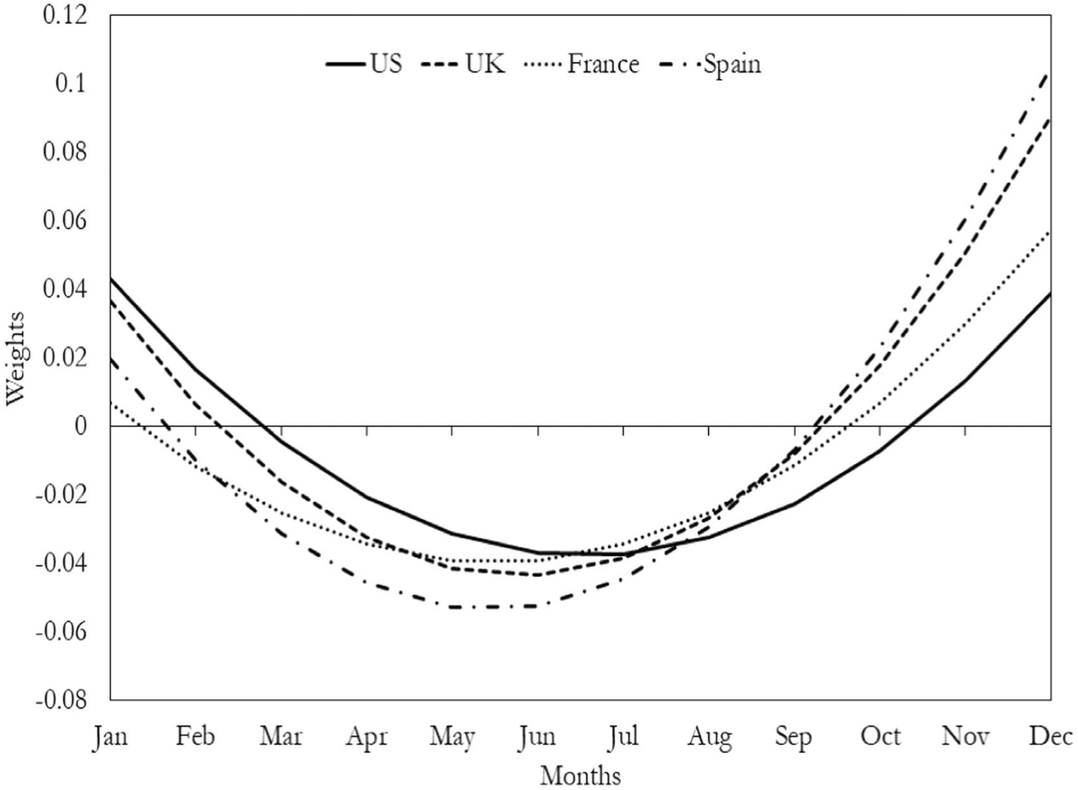

FIGURE 1 Mixed-frequency data sampling (MIDAS) weighting scheme (Almon *P*). This figure plots the MIDAS weighting scheme from regressions of (annual) changes in dividend growth against the lagged (monthly) first difference of the dividend–price ratio. The figure plots the associated weights for a subsample of four countries: United States, United Kingdom, France, and Spain. The MIDAS estimation is based on the Almon lag polynomial of order *P*.

indicate that the impact of the dividend–price ratio on future dividend growth is declining slowly with the number of lags, as opposed to experiencing a more rapid decay.

On the whole, our findings confirm that the trade-off between flexibility and parsimony is not straightforward to evaluate when selecting the optimal weighting scheme in the MIDAS specification. Importantly, the optimal choice appears to vary across countries. Although the simplicity of the Almon *P* polynomial makes it the most commonly selected scheme in our sample, the additional structure imposed by the exponential Almon and the beta polynomials results in a superior fit for a sizable minority of countries.

Regarding the relative importance of different lags, the MIDAS results indicate that the dividends paid toward the end of the year have a substantially greater role when predicting the annual dividend growth, compared to dividends that are paid earlier in the year. For instance, similar to Asimakopoulos et al. (2017), we find that dividends paid in December consistently take the highest weight, whereas dividends paid in the summer tend to take the lowest weights. Figures 1 and 2 illustrate this stylized fact across four sample countries (United States, United Kingdom, 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.

Overall, one significant advantage of the MIDAS approach is that it does not impose strong 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 nonlinear lag polynomials under a MIDAS regression results in a more parsimonious

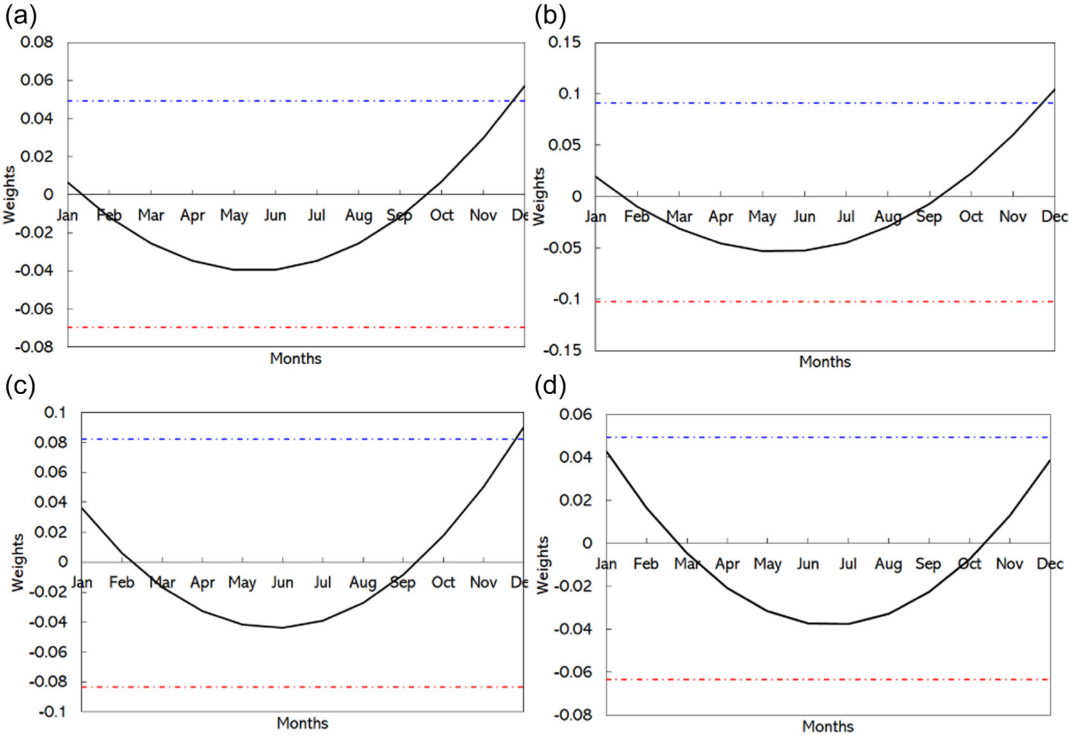

FIGURE 2 Mixed-frequency data sampling (MIDAS) weighting scheme (Almon *P*). This figure plots the MIDAS weighting scheme from regressions of (annual) changes in dividend growth against the lagged (monthly) first difference of the dividend–price ratio. The figure plots the associated weights for a subsample of four countries: (a) France, (b) Spain, (c) United Kingdom, and (d) United States. 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. [Color figure can be viewed at wileyonlinelibrary.com]

estimation with lower sensitivity to specification errors, compared to the alternatives of state-space models or mixed-frequency vector autoregression (VAR) models. Finally, the MIDAS approach avoids parameter proliferation, which could have more pronounced consequences in our sample countries with relatively small data sets.

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 & Shiller, 1988), given by:

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

where ρ denotes an autoregressive parameter. We follow Cochrane (2008) and estimate ρ via a 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 that allows us to obtain ex post rational prices by using the dividend yield and expected dividend growth.

As discussed earlier, although 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 detrend the price and dividend series as $\lambda = e^{\beta}$, where β refers to the slope coefficient from the regression in Equation (2). The real detrended time series of prices and dividends are then given by $\hat{p}_t = p_t/\lambda^{t-T}$ and $\hat{d}_t = d_t/\lambda^{t-T}$, respectively.

Finally, similarly to Shiller (1981), we compute the detrended ex post rational price $\hat{p}_t^*$ recursively from the terminal date T using the following equation:

$$\hat{p}_t^* = \overline{\gamma}\,(\hat{p}_{t+1}^* + \hat{d}_t), \tag{11}$$

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

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

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

Shiller (1981) and LeRoy and Porter (1981) are the first 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 described in (12), the maximum value of the variance of changes in price for a given variance of dividends is given by:

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

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 that are known either many years before or just before they are paid.[8] Finally, a slightly different version of the variance inequality can be written as:

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

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 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 nonstationary because of the general tendency of firms to smooth dividends over time, whereas Merton (1987) argues that the results in Shiller (1981) are more likely to reflect a rejection of the stationarity assumption than a rejection of market efficiency. Because we also find strong evidence of nonstationarity 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 nonstationarity in dividends, the variance of innovations in the stock price must be lower than the variance of innovations in the corresponding dividend.[9] 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, we show that excess volatility inequality is reversed to

$$\mathrm{var}(\Delta\hat{p}_t) \ge \mathrm{var}(\Delta\hat{p}_t^*), \tag{15}$$

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

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: (1) regressions of changes in dividend growth against the lagged dividend–price ratio at an annual frequency (as in Cochrane, 2008) and (2) MIDAS regressions of changes in annual dividend growth against the lagged first differences of the monthly dividend–price ratio (as in Asimakopoulos et al., 2017). For comparability with the literature, we also report the regression results from estimating the original regression of prices against time (as in Shiller, 1981).

Estimating the standard Shiller-type regressions of log price against time in Equation (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 is found to have the lowest slope (β = 0.03), compared to Romania that has only 19 years of data and the highest slope

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

[9]West (1988) computes innovations (i.e., unexpected changes) in the stock price and the corresponding dividend series as the residuals from fitting autoregressive integrated moving average (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) addresses small-sample bias.

TABLE 4 Predictive regressions for dividend growth.

		Shiller-type regressions		Cochrane-type regressions		MIDAS regressions	
Country	*n*	Slope	R^2	Slope	R^2	Slope	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

(Continues)

TABLE 4 (Continued)

Country	n	Shiller-type regressions Slope	R^2	Cochrane-type regressions Slope	R^2	MIDAS regressions Slope	R^2
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

Note: This table reports the estimated intercept and slope coefficients, and the R^2 from predictive least squares regressions for dividend growth. The first se of columns reports the results from regressing stock prices against time (as in Shiller, 1981). The second set of columns reports the results from regressing dividend growth against the lagged dividend–price ratio at an annual frequency (as in Cochrane, 2008). The third set of columns reports the results from mixed-frequency data sampling (MIDAS) regressions of dividend growth against the lagged first difference of the dividend–price ratio. Countries are sorted in ascending order based on the number of years in the second column with available data (n).
*$p < 0.10$; **$p < 0.05$; ***$p < 0.01$.

(β = 0.64). This relation 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 of 50 countries are found to have a positive $\hat{\beta}_{1,g}$, with 11 of these cases being statistically significant at the 5% level. By comparison, of the 15 negative slopes, only 3 are statistically significant at the 5% level (Austria, Finland, and Japan). When we account for intrayear seasonality by estimating the MIDAS specification in (3), the resulting slope coefficients are now universally positive, with 15 of 50 positive slopes also being statistically significant at the 5% level.

In addition, Table 5 reports descriptives statistics for the R^2 obtained from estimating each regression across the 50 countries. In terms of in-sample predictability, accounting for seasonality by using a mixed-frequency regression 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 than that of the Cochrane-type regressions in (2), with a mean R^2 of 55% in the former compared to only

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

Note: The table reports descriptive statistics for the R^2 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 first difference of the 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^2 across the 50 sample countries.

4% in the latter. Furthermore, the R^2 in the MIDAS regressions ranges from 9% to 96% across the 50 countries, whereas the respective range of the R^2 in the Cochrane-type regressions is 0% to 20%. Importantly, the MIDAS specification results in a higher R^2 than the one obtained from the Cochrane-type specification in all 50 countries. Based on this difference in predictive power, we expect the detrending 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.[10]

5.2 | First-generation bounds tests

To get a preliminary idea about the magnitude of the excess volatility puzzle across markets, we begin by plotting in Figure 3 the historical evolution of the detrended stock price $\hat{p}_t$ and the corresponding ex post rational price $\hat{p}_t^*$ for France, Spain, United States, and United Kingdom. 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 $\hat{p}_t^*$ consistently plotting as a much smoother and more stable series compared to that of its respective $\hat{p}_t$. In this sense, Figure 3 provides preliminary evidence against the rational expectations hypothesis in an international context, consistent with the respective figures for the United States in Shiller (1981), which depict the time series of ex post rational prices as considerably smoother than the time series of realized prices.

Figure 4 plots the ratio $\theta = \sigma(\hat{p}_t)/\sigma(\hat{p}_t^*)$ 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 countries. In the interest of comparability with the 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).[11]

[10]Interestingly, regressing stock prices against time (as in Shiller, 1981) also results in high in-sample fit, with a mean R^2 of 53%. Nevertheless, these R^2 values are not directly comparable to those obtained when estimating the predictive regressions in (2) and (3), as 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.

[11]Tables D1–D5 in the Online Appendix present more detailed results of the first-generation Shiller (1981) bounds tests. For instance, we find that the magnitude of excess volatility is 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 United States. For robustness, we also perform those tests when expected dividend growth is 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).

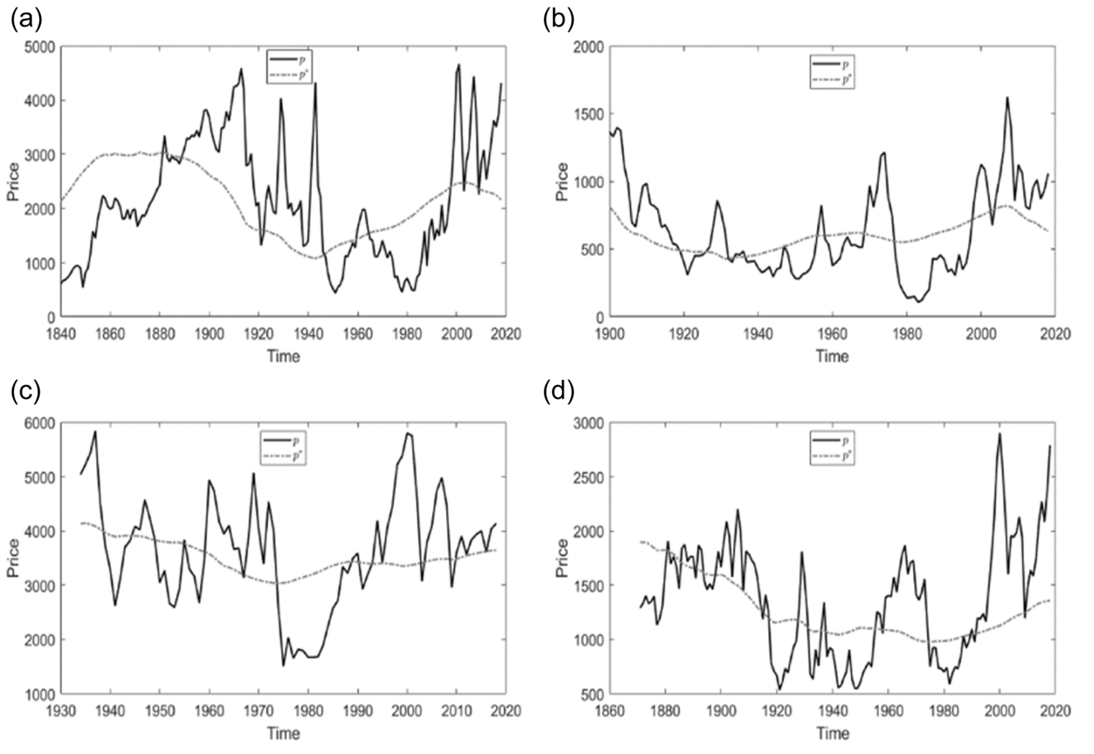

FIGURE 3 Time evolution of realized price and ex post rational price. This figure plots the time series of detrended realized prices ($\hat{p}_t$) and detrended ex post rational prices ($\hat{p}_t$) for a subsample of four countries: (a) France, (b) Spain, (c) United Kingdom, and (d) United States.

The first thing to notice is that $\sigma(\hat{p}_t)$ is indeed higher than the volatility $\sigma(\hat{p}_t^*)$ that would have been expected conditional on dividends across all 50 countries (Spain is the only exception, 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 United States 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 reported in Shiller (1981). For instance, the volatility ratio θ ranges from 0.5 (Spain) to 22.7 (Malaysia), with a mean (median) value of 7.5 (6.0). Overall, in addition to the ratio almost universally exceeding 1 (i.e., $\sigma(\hat{p}_t)$ exceeding $\sigma(\hat{p}_t^*)$), the volatility ratio exceeds 5 (reported in the United States by Shiller, 1981) for 27 of 50 countries.

5.3 | Second-generation bounds tests

The dividends' order of integration may 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, p. 951), in particular, states, *"Given the near-random-walk behavior of stock prices, the volatility of the stock price is captured by* $\text{var}(p_t - p_{t-1})$, *"* in which case the relatively high volatility of the realized price would be consistent with the inequality in (15).

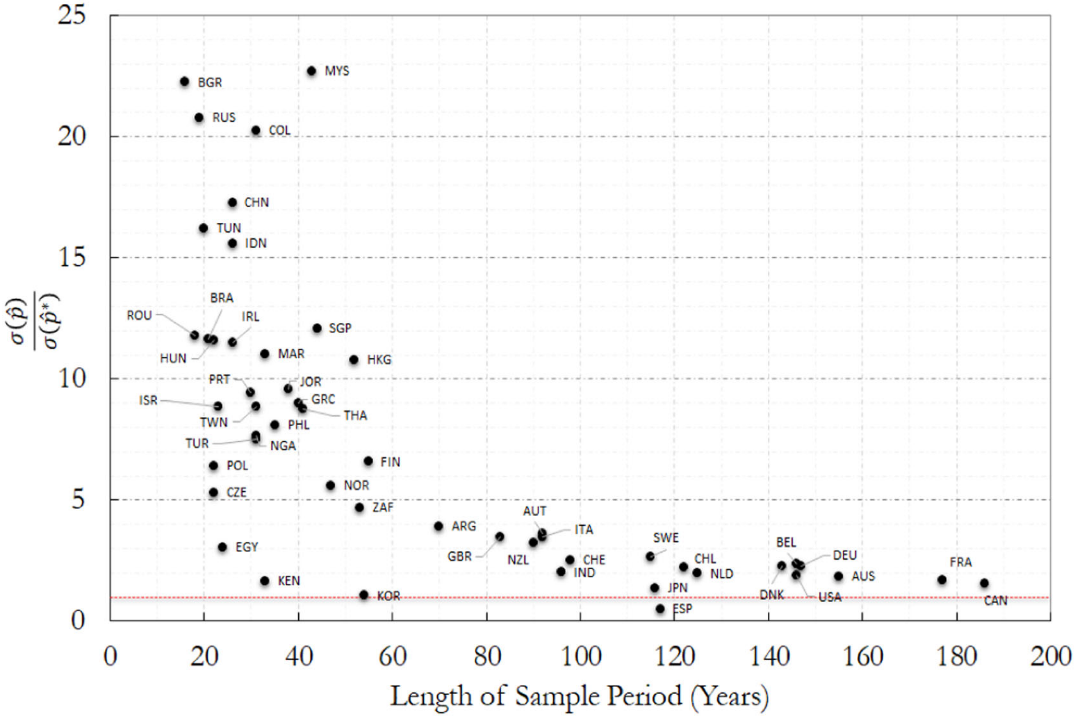

FIGURE 4 Shiller volatility ratios and sample period length. This figure plots the Shiller (1981) ratio $\sigma(\hat{p}_t)/\sigma(\hat{p}_t^*)$ of the volatility of the realized price over the volatility of the ex post rational price across a sample of 50 countries. The volatility ratios are ordered (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 1 violate the Shiller (1981) variance bound.

Because dividends appear 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, given in inequality (15). Table 6 reports a set of summary statistics of the volatility ratio $\sigma(\Delta\hat{p}_t)/\sigma(\Delta\hat{p}_t^*)$ across the 50 countries. Panel A reports the results when expected dividend growth is obtained via the Cochrane-type regressions at an annual frequency in (2), and Panel B reports the results when expected dividend growth is obtained via the MIDAS regressions in (3).[12]

The results fail to reject the market efficiency hypothesis when we account for dividends' nonstationarity via the second-generation Engel (2005) variance bounds test. Irrespective of the measure of expected dividend growth used, the volatility ratio exceeds unity across all sample countries; the Czech Republic is the only exception (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 be expected if markets are efficient.

Importantly, the resulting volatility ratios $\sigma(\Delta\hat{p}_t)/\sigma(\Delta\hat{p}_t^*)$ are markedly different depending on the way in which expected dividend growth is obtained. When we use regressions at an annual frequency (as in

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

TABLE 6 Variance bounds tests: Engel volatility ratio.

	Full sample	Subsamples formed on time-series length		
		L ≤ 30	30 < *L* ≤ 60	60 < *L*
Panel A: Expected growth via Cochrane-type regressions				
Mean	14.1	15.3	15	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				
Mean	2.8	2.3	3.6	2.2
Median	2.2	2.4	2.8	1.7
Min.	1	1	1.2	1
Max.	14.5	4.4	14.5	9
n > 1	49	12	19	18
n	50	13	19	18

Note: This table presents the results of computing the Engel (2005) volatility ratio $\boldsymbol{\sigma}(\boldsymbol{\Delta\hat{p}_t})/\boldsymbol{\sigma}(\boldsymbol{\Delta\hat{p}_t^*})$. 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 (as in Cochrane, 2008). Panel B reports the results when expected dividend growth is obtained via mixed-frequency data sampling (MIDAS) regressions of changes in (annual) dividend growth against the lagged (monthly) first difference of the 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, and the last three columns report the results from subsamples that have been formed based on the countries' number of years of available data (***L***).

Cochrane, 2008), the mean volatility ratio is 14.1, and it ranges from 4.7 (Romania) to 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 2.8, and it ranges from 1.0 (Czech Republic) to 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 those obtained under Cochrane-type regressions at an annual frequency.

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 $\sigma(\Delta\hat{p}_t)$ being substantially closer to $\sigma(\Delta\hat{p}_t^*)$, 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 be expected based on subsequent dividends. This large impact of the mixed-frequency estimation on 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.

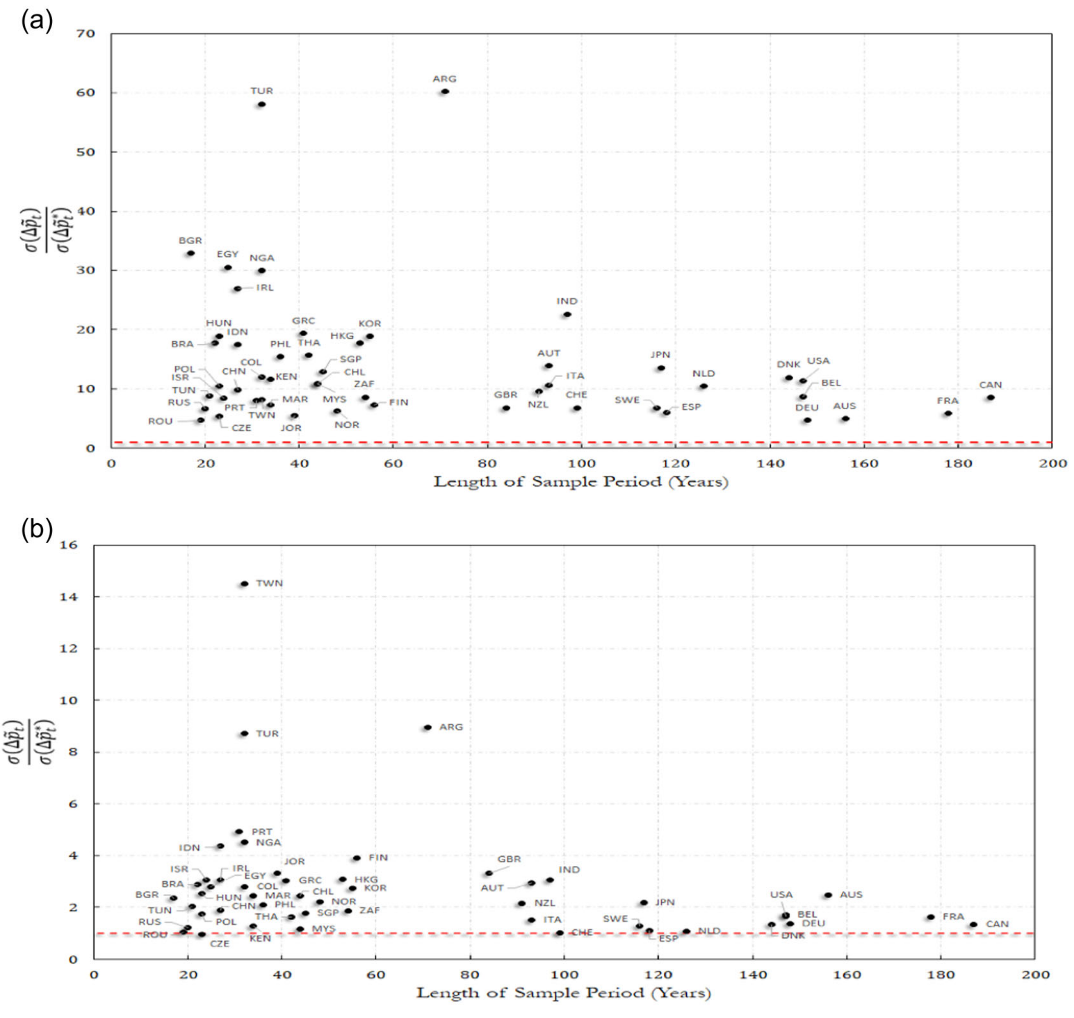

FIGURE 5 Engel volatility ratios and sample period length. This figure plots the Engel (2005) ratio $\sigma(\hat{p}_t)/\sigma(\hat{p}_t^*)$ of the volatility of realized price changes over the volatility of ex post rational price changes across a sample of 50 countries. The volatility ratios are ordered (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 1 violate the Engel (2005) variance bound. Figure 1a 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 (as in Cochrane, 2008), and Figure 1b plots the respective ratios when expected dividend growth is obtained via mixed-frequency data sampling (MIDAS) regressions of changes in (annual) dividend growth against the lagged (monthly) first difference of the dividend–price ratio. [Color figure can be viewed at wileyonlinelibrary.com]

Another interesting finding refers to an observed negative relation 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 in Panel A of Table 6, countries with longer time series of available data tend to have substantially lower volatility ratios compared to countries with relatively shorter time series of available data, although this relation 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, whereas 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 12.4 in the subsample of countries with more than 60 years of available data, whereas the respective mean is 15.3 for countries with fewer than 30 years of available data.[13]

In contrast, this negative relation between sample length and 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 ratios are 2.3 for countries with fewer than 30 years of available data and 2.2 for countries with more than 60 years of data.

5.4 | Robustness

We consider a set of additional tests to determine the robustness of our main empirical findings regarding the importance of accounting for dividend seasonality via MIDAS regressions. The results of these robustness checks are briefly summarized below and reported in the Online Appendix.

1. *Detrending factor*: In the interest of comparability with the literature, we repeat the empirical analysis using the detrending factor obtained from Shiller (1981) regressions of prices against time, ignoring the issue of nonstationarity of prices. The results are reported in Tables D1–D5 and Figure D1 in the Online Appendix.
2. *MIDAS scheme selection*: We perform the forecasting exercise for selecting the optimal MIDAS weighting scheme by adopting the alternative 70/30 rule. In this setting, we use the first 70% of the full sample period as the estimation period and the remaining 30% of the sample period for the out-of-sample evaluation of forecasts obtained by competing schemes. The results are broadly similar to those reported in Table 3, with the Almon *P* polynomial being the most often selected scheme. The performance of alternative weighting schemes under the 70/30 rule is reported in Tables H1 and H2 in the Online Appendix.

   We also compute Engel (2005) volatility ratios when the MIDAS scheme is selected using the 70/30 rule, as opposed to a fixed 2-year out-of-sample period. The results are qualitatively the same as those reported in Table 6, with MIDAS volatility ratios being substantially lower than those obtained under Cochrane-style regressions of dividend growth. These results are reported in Table H3 in the Online Appendix.
3. *Quarterly aggregation*: We compute the Engel (2005) volatility ratio using quarterly (instead of monthly) dividends in the MIDAS specification. The results are qualitatively the same as those reported in Panel C of Table 6, with the resulting volatility ratios being substantially lower than those obtained using Cochrane-type regressions with annual dividends. These results are reported in Table G1 and Figures G1–G3 in the Online Appendix.
4. *World War II (WWII)*: Considering that dividends are commonly found to behave differently before and after WWII (e.g., Chen et al., 2012), we repeat our empirical analysis separately for the pre- and post-WWII subperiods. The results are reported in Tables I1–I6 in the Online Appendix. Overall, our main findings hold for both subperiods. We find that the null of stationarity is consistently rejected before and after WWII (Tables I1–I3). More important, accounting for dividend seasonality significantly improves dividend predictability (Table I5) and results in substantially lower volatility ratios compared to those obtained via Cochrane-type regressions at an annual frequency (Table I6), consistent with our findings in the full sample.

[13]The negative relation between the volatility ratio and 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, whereas countries with more than 60 years of data have a mean ratio of 2.4. Tables D1–D5 in the Online Appendix report more detailed results about the relation between excess volatility and sample length for first-generation bounds tests.

6 | TRADING STRATEGY

Our empirical results thus far are consistent with significant violations of the conventional variance bounds across multiple international markets. We now investigate whether this excess volatility of realized prices relative to expected prices constitutes evidence to reject the hypothesis of weak form rationality in terms of allowing for exploitable profit opportunities.

Bulkley and Tonks (1992) argue that some part of realized prices' excess volatility can be attributed to revisions in the parameters of the model that market participants use for the dividend process. In this setting, investors are rational and use unbiased techniques when forming their expectations about future dividend growth. However, the fact that investors dynamically revise the parameters of their unbiased model results in some excess volatility of realized prices relative to that of expected prices based on future discounted dividends. Motivated by this argument, we explore whether the relatively high volatility of p could be attributed to market participants revising the parameters of their unbiased structural model for dividends. To this end, we construct a simple "buy low/sell high" trading strategy that uses currently available data to set the parameters of the trading rule.[14]

Our strategy is based on the trading rule adopted in Bulkley and Tonks (1992). More specifically, we consider an investor who can choose to invest either in the stock index or in a risk-free bond. At any point in time t, the investor will decide how to invest by comparing the actual index price p against the perfect-foresight price p^*. If the realized price of the index is at least Δ% higher compared to the expected price, the investor will choose to sell the index and buy the risk-free bond. The bond is then held until the realized price falls at least Δ% below the expected price, at which point the investor will buy back into the index. The investor receives dividend payments when holding the index and the risk-free interest when holding the bond, reinvesting all returns.

Under this simple trading rule, the investor will sell overpriced shares when their price exceeds the present value of future dividends (i.e., when $p > p^* \times (1 + \Delta)$) and buy them back when they become underpriced in terms of their price falling below the present value of expected dividends (i.e., when $p < p^* \times (1 + \Delta)$). If markets are efficient, deviations of realized prices from their expected values would not be large enough for this trading rule to dominate a simple buy-and-hold strategy, especially considering the unavoidable model error around expected prices and the low rate of return of the risk-free asset (consistent with the well-documented equity premium puzzle discussed in Mehra & Prescott, 1985). Bulkley and Tonks (1992) provide a comprehensive discussion of the rationale behind this trading rule.

The investor starts with an initial wealth of 100 across any sample country, and the decision to switch in or out of the index is made on December 31 of each year. We adopt a dynamic threshold Δ_t for the investor's choice to switch between the index and the risk-free bond. In particular, at each time t we select a value for Δ_t that would ex post maximize the investor's profits from the trading strategy during the period $(0, t - 1)$.

Table 7 presents the performance of this trading strategy against that of the simple buy-and-hold benchmark, separately for each of the 50 countries. More specifically, we report each strategy's annual volatility, mean annual return, and 5% value-at-risk (VaR) (computed parametrically).[15] In addition to these measures, we report the p-value from the Linton et al. (2005) second-order stochastic dominance (SSD) test of the trading strategy against the buy-and-hold benchmark. The null hypothesis of the SSD test is that the trading strategy TS stochastically dominates the buy-and-hold BH (i.e., $H0 : TS \succ_2 BH$), which represents the case where every risk-averse investor prefers TS over BH, irrespective of the specifics of their risk preferences. Therefore, a rejection of the null hypothesis constitutes evidence in favor of market efficiency because the trading strategy based on excess volatility would not lead to

[14]Shiller (1981) and Bulkley and Tonks (1992) discuss cases where large excess volatility of stock returns suggests the existence of trading rules that may dominate the standard buy-and-hold strategy.

[15]Expected prices p^* are computed via the MIDAS regressions approach. For robustness, we replicated this exercise with expected prices that were obtained via Shiller-type and Cochrane-type regressions. The results are qualitatively similar and thus are not reported for brevity, but they are available upon request from the authors.

TABLE 7 Trading strategy performance.

Country	Volatility		Return		VaR		SSD test
	TS	*BH*	*TS*	*BH*	*TS*	*BH*	$TS \succ_2 BH$
Argentina	1.41	6.72	0.34	1.82	−0.37	−0.83	0.00
Australia	0.09	0.14	0.05	0.05	−0.09	−0.18	0.78
Austria	0.02	0.32	0.06	0.09	0.03	−0.38	0.77
Belgium	0.06	0.22	0.06	0.05	−0.04	−0.29	0.79
Brazil	0.08	0.34	0.18	0.21	0.06	−0.31	0.73
Bulgaria	0.18	0.44	0.09	0.23	−0.14	−0.76	0.51
Canada	0.11	0.17	0.05	0.05	−0.14	−0.23	0.79
Chile	0.01	0.20	0.06	0.11	0.04	−0.21	0.42
China	0.01	0.35	0.04	0.08	0.03	−0.57	0.62
Colombia	0.03	0.45	0.09	0.20	0.04	−0.40	0.37
Czech	0.02	0.25	0.03	0.09	0.01	−0.35	0.59
Denmark	0.08	0.19	0.04	0.05	−0.09	−0.23	0.77
Egypt	0.01	0.64	0.16	0.32	0.13	−0.62	0.47
Finland	0.07	0.35	0.09	0.14	−0.03	−0.39	0.62
France	0.09	0.21	0.05	0.07	−0.08	−0.25	0.53
Germany	0.04	0.48	0.05	0.08	−0.02	−0.53	0.77
Greece	0.09	0.42	0.14	0.14	0.00	−0.57	0.72
Hong Kong	0.03	0.27	0.04	0.08	0.00	−0.4	0.69
Hungary	0.02	0.29	0.07	0.11	0.04	−0.39	0.64
India	0.05	0.27	0.06	0.09	−0.03	−0.31	0.72
Indonesia	0.02	0.14	0.08	0.13	0.05	−0.10	0.31
Ireland	0.03	0.25	0.05	0.10	0.01	−0.40	0.51
Israel	0.02	0.27	0.05	0.11	0.02	−0.30	0.37
Italy	0.07	0.35	0.06	0.11	−0.05	−0.38	0.51
Japan	0.18	0.30	0.05	0.07	−0.25	−0.48	0.73
Jordan	0.02	0.27	0.06	0.08	0.03	−0.33	0.78
Kenya	0.03	0.25	0.13	0.05	0.09	−0.40	0.63
Korea	0.09	0.32	0.07	0.15	−0.08	−0.32	0.26
Malaysia	0.21	0.28	0.08	0.10	−0.25	−0.36	0.72
Morocco	0.01	0.21	0.05	0.08	0.03	−0.26	0.45
Netherlands	0.05	0.19	0.04	0.05	−0.04	−0.26	0.77
New Zealand	0.08	0.22	0.06	0.06	−0.06	−0.26	0.75

TABLE 7 (Continued)

	Volatility		Return		VaR		SSD test
Country	***TS***	***BH***	***TS***	***BH***	***TS***	***BH***	***TS*** $\succ_2$ ***BH***
Nigeria	0.15	0.35	0.2	0.25	−0.02	−0.36	0.71
Norway	0.04	0.30	0.07	0.11	0.01	−0.39	0.74
Philippines	0.05	0.28	0.10	0.08	0.01	−0.41	0.70
Poland	0.09	0.31	0.10	0.11	−0.03	−0.41	0.67
Portugal	0.05	0.31	0.08	0.06	0.00	−0.44	0.74
Romania	0.06	0.54	0.09	0.27	0.00	−0.62	0.49
Russia	0.08	1.05	0.11	0.45	0.00	−0.67	0.25
Singapore	0.01	0.27	0.03	0.08	0.01	−0.37	0.52
South Africa	0.14	0.23	0.13	0.15	−0.08	−0.21	0.81
Spain	0.04	0.22	0.06	0.06	0.00	−0.29	0.77
Sweden	0.12	0.23	0.06	0.09	−0.13	−0.28	0.54
Switzerland	0.06	0.20	0.03	0.06	−0.09	−0.28	0.35
Taiwan	0.02	0.29	0.03	0.06	0.00	−0.43	0.68
Thailand	0.25	0.32	0.12	0.12	−0.25	−0.42	0.76
Tunisia	0.01	0.20	0.07	0.12	0.05	−0.17	0.23
Turkey	0.19	1.68	0.38	0.84	0.08	−0.72	0.49
United Kingdom	0.09	0.17	0.07	0.07	−0.07	−0.21	0.80
United States	0.15	0.17	0.05	0.06	−0.20	−0.24	0.54

Note: This table reports the performance of a trading strategy that is based on the difference between realized prices p and expected prices p^*, with the latter estimated via mixed-frequency data sampling (MIDAS) regressions of dividend growth against the lagged dividend–price ratio. The performance of the trading strategy (***TS***) is compared against that of the simple buy and-hold benchmark (*BH*). The table reports annual volatility (in **%**), mean annual return (in **%**), and 5% value-at-risk (**VaR**) quantile. The last column reports the p-value of the Linton et al. (2005) second-order stochastic dominance (**SSD**) test. The null hypothesis of the SSD test is that the trading strategy stochastically dominates the buy-and-hold benchmark (i.e., **H0 : :*TS*** $\succ_2$ ***BH***).

performance that is universally preferred by every investor, with the ordering of *TS* relative to *BH* being instead dependent on individual investors' particular utility functions.

As can be seen in Table 7, the buy-and-hold benchmark offers annual returns that are on average higher (for 46 of 50 countries) and have a higher level of volatility (for all 50 countries) relative to the trading strategy. This is expected, as the trading strategy involves some periods where investors switch from holding risky stocks to risk-free bonds that offer lower returns at a lower volatility level.

However, the trading strategy loads consistently less on downside risk, as evidenced by a VaR quantile to the right of the corresponding VaR quantile of the benchmark, across all 50 countries. There are, in fact, some countries for which the VaR quantile is positive, such as Austria, Brazil, Chile, China, Colombia, Indonesia, and Turkey, suggesting that the trading strategy would have produced mainly positive returns over the corresponding period.

The SSD test fails to reject the null of the trading strategy dominating the buy-and-hold benchmark in almost every sample country (except for Argentina). These results suggest that any risk-averse investor in almost every country would prefer trading based on the difference between realized prices and expected prices, compared to simply buying and holding the index, irrespective of their utility function. In other words, the excess volatility of realized prices consistently results in profitable investment opportunities that dominate the buy-and-hold strategy. Therefore, investors' revision of the parameters of their unbiased model for dividends is an unlikely explanation for the magnitude of excess volatility that is observed across multiple countries, given the extent of the trading strategy's stochastic dominance over the benchmark. Overall, the SSD test provides strong evidence against the efficient markets hypothesis, consistent with the arguments in Shiller (1981) and Bulkley and Tonks (1992).

7 | 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 finding excess volatility across all our sample countries. However, given that in reality dividends deviate substantially from stationarity, these results seem to say more about the assumption's validity than about market efficiency. Second-generation bounds tests, which allow for nonstationary dividends, fail to reject the efficient market hypothesis in any of the sample countries, casting doubt on the existence of the excess volatility puzzle.

ORCID

Nikolaos Voukelatos http://orcid.org/0000-0001-8272-2901

REFERENCES

Almon, S. (1965). The distributed lag between capital appropriations and expenditures. *Econometrica*, *33*, 178–196.

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

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

Asgharian, H., Hou, A. J., & 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., & 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., & Wright, J. H. (2013). State space models and MIDAS regressions. *Econometric Reviews*, *32*(7), 779–813.

Breitung, J., & Roling, C. (2015). Forecasting inflation rates using daily data: A non-parametric midas approach. *Journal of Forecasting*, *34*, 588–603.

Bulkley, G., & Tonks, I. (1992). Trading rules and excess volatility. *Journal of Financial and Quantitative Analysis*, *99*, 1083–1098.

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

Chen, L., Da, Z., & Priestley, R. (2012). Dividend smoothing and predictability. *Management Science*, *58*, 1–20.

Clements, M. P., & 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., & Ghysels, E. (2007). Why do absolute returns predict volatility so well? *Journal of Financial Econometrics*, *5*(1), 31–67.

Gagliardini, P., Ghysels, E., & 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., & Qian, H. (2019). Estimating MIDAS regressions via OLS with polynomial parameter profiling. *Econometrics and Statistics*, *9*, 1–16.

Ghysels, E., Santa-Clara, P., & 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., & 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., & Valkanov, R. (2007). MIDAS regressions: further results and new directions. *Econometric Reviews*, *26*(1), 53–90.

Golez, B., & 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., & LeRoy, S. (2014). Risk aversion, investor information and stock market volatility. *European Economic Review*, *70*, 88–107.

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

Linton, O., Maasoumi, E., & Whang, Y.-J. (2005). Consistent testing for stochastic dominance under general sampling schemes. *Review of Economic Studies*, *72*(3), 735–765.

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

Mehra, R., & Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, *15*(2), 145–161.

Merton, R. (1987). On the current state of the stock market rationality hypothesis. In R. Dornbusch, S. Fischer, & J. Bossons (Eds.), *Macroeconomics and finance: Essays in honor of Franco Modigliani* (pp. 93–124). MIT Press.

Mogliani, M., & Simoni, A. (2021). Bayesian MIDAS penalized regressions: estimation, selection, and prediction. *Journal of Econometrics*, *222*, 833–860.

Pettenuzzo, D., Timmermann, A., & Valkanov, R. (2016). A MIDAS approach to modeling first and second moment dynamics. *Journal of Econometrics*, *193*, 315–334.

Rangvid, J., Schmeling, M., & 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., & 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.

You, Y., & Liu, X. (2020). Forecasting short-run exchange rate volatility with monetary fundamentals: A GARCH-MIDAS approach. *Journal of Banking & Finance*, *116*, 105849. https://doi.org/10.1016/j.jbankfin.2020.105849

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Quaye, E., Tunaru, R., & Voukelatos, N. (2024). MIDAS and dividend growth predictability: Revisiting the excess volatility puzzle. *Journal of Financial Research*, 1–25. https://doi.org/10.1111/jfir.12403