Improving U.S. Stock Return Forecasts: *A "Fair-Value" CAPE Approach*

JOSEPH DAVIS, ROGER ALIAGA-DÍAZ, HARSHDEEP AHLUWALIA, AND RAVI TOLANI

JOSEPH DAVIS
is global chief economist and global head of the Investment Strategy Group at Vanguard in Valley Forge, PA.
joseph_davis@vanguard.com

ROGER ALIAGA-DÍAZ
is chief economist, North America and principal of the Investment Strategy Group at Vanguard in Valley Forge, PA.
roger_aliaga-diaz@vanguard.com

HARSHDEEP AHLUWALIA
is a senior investment strategist at Vanguard in Valley Forge, PA.
harshdeep_singh_ahluwalia@vanguard.com

RAVI TOLANI
is currently a student at the Fuqua School of Business at Duke University in Durham, NC, and was an investment analyst at Vanguard at the time this article was written.
ravigtolani@gmail.com

Professors John Campbell and Robert Shiller's [1988, 1998] cyclically adjusted price/earnings (P/E) ratio (Shiller CAPE ratio) is arguably the most widely followed metric in the investment profession to judge whether or not a stock market is fairly valued. The CAPE ratio's popularity is due in part to the power of mean reversion. A high (low) CAPE ratio has been associated with below-average (above-average) 10-year-ahead U.S. stock returns.

Nevertheless, stock return predictions using the Shiller CAPE ratio have recently performed more poorly. Beginning around 1985, the average out-of-sample forecast errors of the predicted returns 10 years ahead (i.e., 1995 and onward) have been larger than if one had used the trailing historical long-run average. The rise in average forecast error has coincided with the secular rise in the CAPE ratio above its 1926–1984 average of 14.6. The Shiller CAPE ratio has defied mean reversion since that time, having only once dropped below its long-run average. Realized U.S. stock returns over the past three decades have been robust, notwithstanding the global financial crisis.

The combination of elevation in CAPE ratios and its recent deterioration in stock-return forecasts has led to a bit of a renaissance in research refining the CAPE-based forecasting framework that has become standard in the investment community. Some recent studies have focused on refining how the Shiller CAPE ratio is constructed. For instance, Siegel [2016] argued that the secular rise in the CAPE ratio's trend is due to changes in accounting standards and that national income and product account (NIPA) earnings should serve as the "E" in the CAPE ratio. Philips and Ural [2016] evaluated alternative weighting schemes to construct the CAPE ratio, including revenue and gross domestic product (GDP) weights and sector composition. Collectively, the improvements in forecast accuracy using these alternative approaches are somewhat mixed, at least for the U.S. stock market.

The economic environment has been cited as another factor in (justifying) elevated CAPE ratios. On a keynote panel at the 70th Annual CFA National Conference in May 2017, Professors Jeremy Siegel and Robert Shiller both cited low interest rates as a potential factor in the extended period of elevated CAPE ratios, although neither explicitly quantified the link between interest rates and future stock returns. This article does just that.

Here we break the standard assumption that the CAPE ratio will revert mechanically to its fixed long-run average, regardless of the economic environment. We disagree with Philips and Ural [2016] that the CAPE ratio does not have a steady-state level. Rather, we stipulate that the steady-state or fair-value CAPE ratio (i.e., the value to which the actual CAPE ratio should eventually revert) varies over time, depending on the state of the economy, as measured by real interest rates, expected inflation, and measures of financial volatility. In our framework, lower real bond yields imply lower real earnings yields and a higher fair-value CAPE ratio, all else equal. Real yields matter in our framework, not nominal yields per se as in the so-called Fed model (Asness [2003]).

Our methodology is most similar to the pioneering work that Bogle [1991, 1995] and Bogle and Nolan [2015] published previously in *The Journal of Portfolio Management*. The so-called Occam's razor model of Bogle and Nolan [2015] projects 10-year-ahead U.S. stock returns based on the current level of the dividend yield, the trailing 10-year average in earnings growth, and a straight-lined reversion of the current P/E ratio to its trailing 30-year average. We attempt to maintain the elegant simplicity of Bogle and Nolan's approach, while refining and improving upon the assumption of—and economic rationale for—CAPE mean reversion. Both approaches tend to produce similar stock forecasts *unless* real bond yields differ from their long-run average at the time that the stock market forecast is made. That is certainly the case today; as of December 31, 2016, our derived real 10-year Treasury yield was near 0%.

Our model's out-of-sample forecasts outperform the conventional approaches using Shiller's CAPE ratio, Siegel's CAPE ratio, and even the Occam's razor model of Bogle and Nolan [2015]. Real-time forecast errors for 10-year-ahead U.S. stock returns have been lower since 1960, and significantly lower since 1985. Specifically, the average return forecast error of our two-step approach since 1985 is 4.1% (3.4%), versus 7.8% (6.2%) from a linear predictive regression using the Shiller (Siegel) CAPE ratio.

We conclude with a discussion of our model's low U.S. equity return projections over the next decade through 2026. In short, low real bond yields justify higher CAPE ratios today versus historical averages, yet they are very likely to prove insufficient in generating average stock returns over the next decade.

THE CAPE RATIO'S CONVENTIONAL USAGE

Future long-run U.S. stock returns have tended to move inversely with the CAPE ratio over time.[1] Typically, financial analysts express monthly annualized 10-year-ahead stock returns as a linear function of the latest Shiller CAPE ratio via an ordinary least squares predictive regression:

$$R_{t+120} = \alpha + \beta \text{CAPE}_t + \epsilon_t \qquad (1)$$

The CAPE ratio has explained a strikingly high 54% (41%) of the time-series variation in 10-year-ahead nominal (real) U.S. stock returns, as measured by Equation (1)'s in-sample, or fitted, R^2, over the 1926–2016 period. Further enhancing the popularity of Shiller's CAPE ratio is that it peaked in 1929 and 1999, before noted stock market crashes.

CAPE RATIO'S FORECAST ACCURACY HAS DETERIORATED LATELY

Unfortunately, the CAPE ratio's out-of-sample forecast accuracy has weakened since the mid-1980s versus its in-sample fit, as illustrated in Exhibit 1. To be sure, the correlations between actual U.S. stock returns and those predicted 10 years prior by the Shiller CAPE ratio have remained high in real time. The correlation has been 83% since 1960 and a remarkable 91% since 1985. But there is an important catch.

We must stress that correlation is not necessarily a reliable indicator of forecast accuracy.

A better measure of forecast accuracy is the average forecast error (i.e., RMSE) between the actual and predicted rolling 10-year-ahead returns. To stress this distinction, Exhibit 2 presents rolling actual versus the predicted 10-year annualized U.S. stock returns since 1960. Whereas the CAPE-based predictions using Equation (1) have been highly correlated with the actual future returns, the forecast error in absolute returns (the basis for RMSE) has generally grown over time. Beginning with long-run forecasts made in the mid-to-late 1980s, the Shiller CAPE ratio's projected stock returns have generally been too bearish, even when one includes the 1999 tech bubble. Put another way, real-time investors would have been better served by using the historical average return as the baseline forecast of future stock returns over the past two decades.

The changing composition of the U.S. stock market is unlikely to be the driving factor here given the high and tight correlation among CAPE ratios weighted

Exhibit 1

The CAPE Ratio's Predictive Power Out-of-Sample

Predictive Variable	Out-of-Sample Forecasts Made Since 1960: Correlation of Predicted Returns with Actual	Average Forecast Error (RMSE)	Model Forecast Error Relative to Error of Using a Naïve Historical Mean Forecast	Out-of-Sample Forecasts Made Since 1985: Correlation of Predicted Returns with Actual	Average Forecast Error (RMSE)	Model Forecast Error Relative to Error of Using a Naïve Historical Mean Forecast
Panel A: Nominal Returns						
Historical average		5.8%			6.2%	
Shiller CAPE ratio	83%	5.5%*	LOWER	91%	7.8%***	HIGHER
Siegel CAPE ratio	67%	4.9%***	LOWER	90%	6.2%	SIMILAR
Panel B: Real Returns						
Historical average		6.4%			5.7%	
Shiller CAPE ratio	56%	6.3%	SIMILAR	81%	7.8%***	HIGHER
Siegel CAPE ratio	36%	6.2%	SIMILAR	76%	5.8%	SIMILAR

Notes: The statistics shown are for 10-year annualized returns using the traditional predictive regression from Equation (1) with Shiller CAPE and Siegel CAPE. Asterisks next to the root-mean square error (RMSE) refer to the significance of the Diebold–Mariano test [2002] of whether the forecast is statistically better or worse than the historical mean.

Significance levels at 90%, 95%, and 99% are denoted by one, two, and three asterisks, respectively.

Source: Authors' calculations.

Exhibit 2

The CAPE Ratio's Real-Time Forecasts since 1960

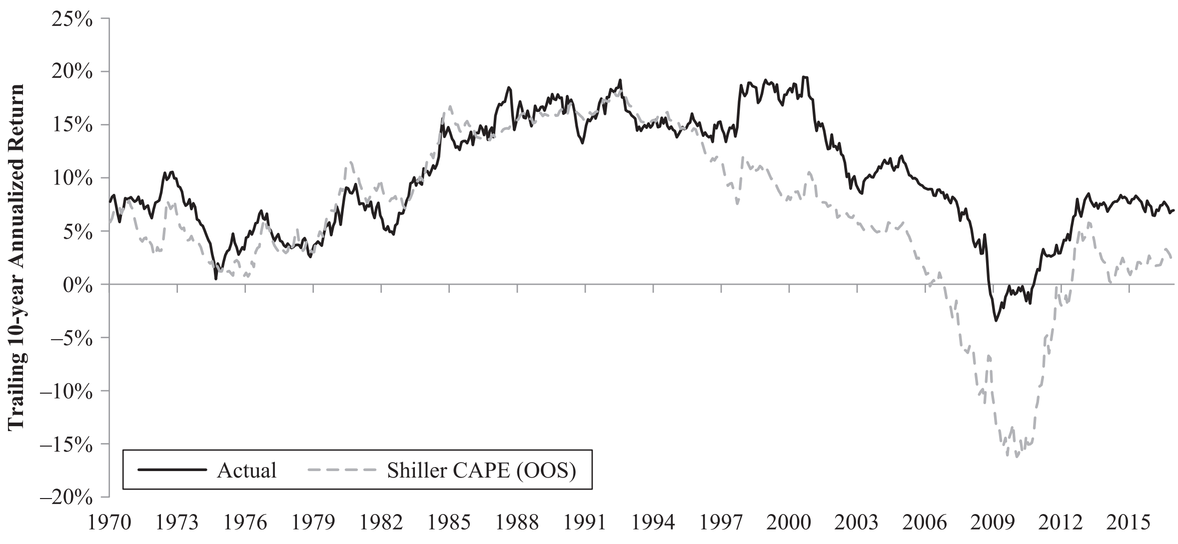

Notes: For the real-time analysis, the regression coefficients are determined recursively at a monthly frequency, starting with January 1926–December 1959 data and re-estimating the regression coefficients every month thereafter. The gap between the two lines represents forecast error.

Source: Authors' calculations.

by the U.S. Consumer Price Index (CPI), GDP, and revenues, respectively, as demonstrated by Philips and Ural [2016, especially Exhibit 2]. Statistically, the CAPE ratio's degradation in real-time forecasting power is due to the lack of mean reversion in the CAPE ratio itself. Exhibit 3 reveals that the CAPE ratio has drifted secularly upward since the late 1980s. Indeed, it has only dropped below its long-run 1926–2016 mean once since

EXHIBIT 3

To Which Mean Will the CAPE Ratio Revert?

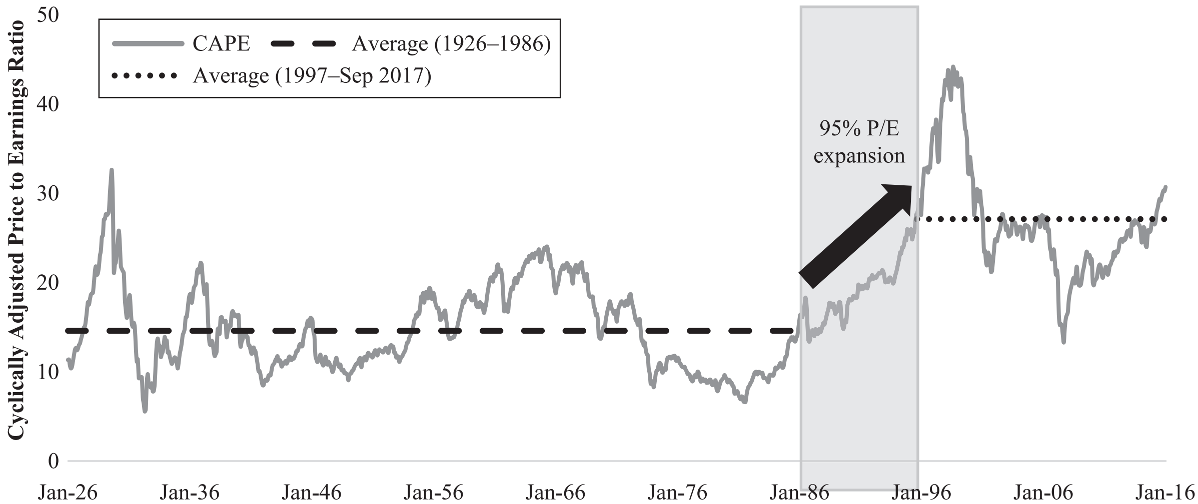

Source: Calculations based on the data obtained from Robert Shiller website, at aida.wss.yale.edu/~shiller/data.htm.

that time, albeit briefly, during the global financial crisis of 2009. There could be several reasons for this.

Siegel [2016] argued the rise in the CAPE's trend is primarily due to changes in accounting standards over time and that NIPA earnings should be substituted for generally accepted accounting principles earnings when applying the CAPE ratio. The bottom row of Exhibit 1 shows that although real-time return projections since 1960 are marginally better using Siegel's CAPE ratio, its forecasts since 1985 have been statistically equivalent to the historical average, having roughly the same RMSE. Regardless of how we define or smooth earnings here, the real-time forecasting accuracy has been weaker than its in-sample fit. Changing the definition of "E" in the CAPE ratio may help, but it does not appear to be a panacea on its own for forecasting U.S. stock returns.

THE ISSUE IS NOT WITH THE CAPE RATIO, BUT WITH CAPE REGRESSIONS

The weak predictability of the CAPE ratio may be less about the earnings and weights used in its calculation and more a reflection of *model instability* (Goyal and Welch [2008]; Pettenuzzo and Timmermann [2011]). In other words, the estimated parameters in Equation (1) for the average return to which stocks revert (dictated by the regression's conditional mean, $\propto$) and the speed of the convergence (as governed by the regression's slope, β) have not remained constant over time.

As evidence that traditional CAPE regressions suffer from model instability, Exhibit 4 presents the results of cumulative sum (CUSUM) tests of Equation (1) using the Shiller and Siegel CAPE ratios, respectively. The lines of the CUSUM test signify parameter instability of conventional CAPE regressions, as the solid line breaches the 5% significance lines around 1985 or so. Exhibit 4 helps to explain the weak out-of-sample performance we document for both Shiller's and Siegel's CAPE ratios in Exhibit 1 despite the high in-sample correlation estimated by Equation (1).

MEAN REVERSION IS CONDITIONAL ON THE ECONOMY

CAPE regression instability originates from at least two sources. The first is the estimation bias that arises when persistent (or slow moving) variables such as the CAPE ratio are used to forecast long-run returns (Stambaugh [1999]). The second relates to standard CAPE regressions omitting the explicit relationship between the expected return on equity (i.e., the real earnings yield) and the expected real discount rate or cost of capital (i.e., real bond yields). If changes in long-term real interest rates influence the steady-state or fair-value

E X H I B I T   4

Traditional CAPE Regressions Are Unstable

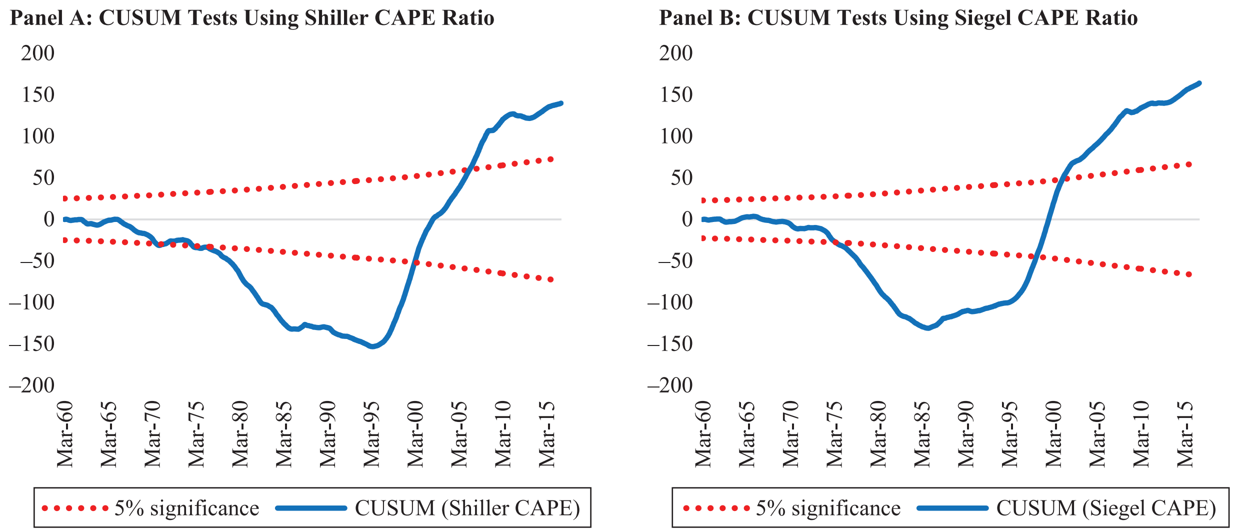

Notes: The CUSUM test (Brown, Durbin, and Evans [1975]) for the 10-year-ahead stock return regression is based on the cumulative sum of the recursive residuals. The test finds parameter instability if the cumulative sum (solid line) extends beyond the area between the two dashed 5% significance lines.

Sources: Authors' calculations, based on data from Robert Shiller website, at aida.wss.yale.edu/~shiller/data.htm; U.S. Bureau of Labor Statistics, and Federal Reserve Board.

CAPE ratio to which stock returns should revert, then the coefficients in traditional CAPE regressions will suffer from instability whenever there are meaningful changes in the level of real bond yields.

This is precisely what we find. The solid lines in Exhibit 4 identify two major periods of instability for the traditional CAPE regression in Equation (1): the late 1970s to mid-1990s and the post-mid-2000s. This parameter instability implies that the CAPE ratio (and its inverse 1/CAPE, or real earnings yield) may not revert mechanically to a fixed, average mean. The low *real* interest rate environment may also explain why the CAPE ratio has not dropped below its long-run average of 16.9 since 1990, except for a brief time during the global financial crisis of 2009. The parameter instability in the CAPE regression appears to coincide with material shifts in average real bond yields, such as the high average real yields between the late 1970s and mid-1990s and the secularly lower real yields before and after that period.

AN IMPROVED TWO-STEP APPROACH USING THE CAPE RATIO

Our hypothesis is simple: Lower real bond yields should imply lower earnings yields and thus higher equilibrium or fair-value CAPE ratios, all else equal. The correlation between real bond and earnings yields in Exhibit 5 suggests that this may be a reasonable approach.[2]

Motivated by this relationship, we propose a two-step approach to forecast stock returns. Although our model differs from the approach typically taken in Equation (1), it can be estimated in real-time using standard software; it does not involve look-ahead bias; and, for the U.S. stock market, it only requires the variables in the CAPE ratio data file conveniently provided by Professor Robert Shiller's website.

Our methodology is most similar to the original work of Bogle [1991, 1995] and Bogle and Nolan [2015]. The so-called Occam's razor model of Bogle and Nolan [2015] projects 10-year-ahead U.S. stock returns based on the current level of the dividend yield, the trailing 10-year-average in earnings growth, and a straight-lined reversion of the current P/E ratio to its trailing 30-year average. We attempt to maintain the elegant simplicity of Bogle and Nolan's approach while refining and improving upon the assumption of—and economic rationale for—CAPE mean reversion. Both approaches should produce similar stock forecasts unless real bond

EXHIBIT 5

The Intuition: Higher Real Bond Yields = High-Equity Earnings Yields

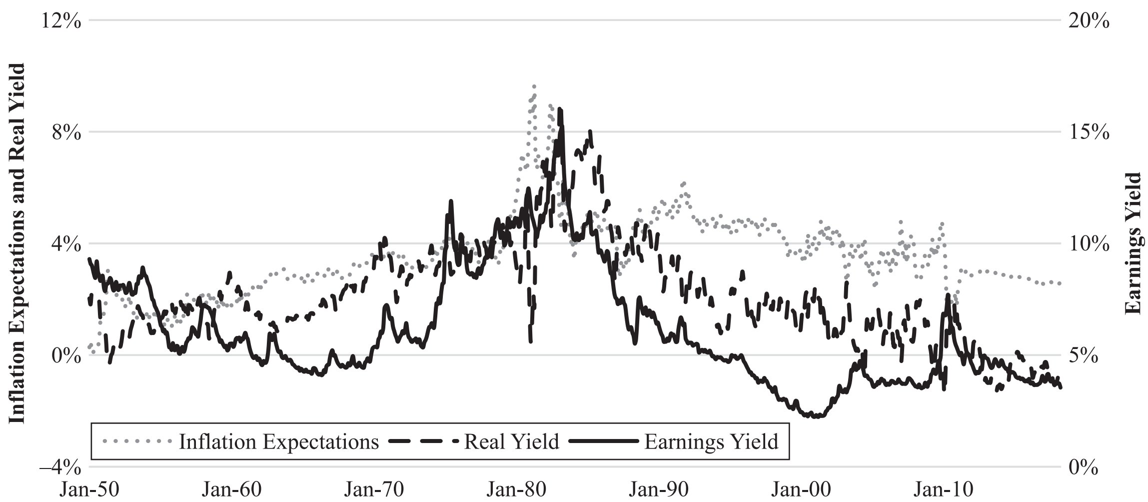

Source: Authors' calculations. Please see the Appendix.

yields differ from their long-run average at the time that the stock market forecast is made.

Step 1: A Vector Autoregressive Model with Earnings Yields, 1/CAPE

Unlike traditional methods, we do not forecast returns directly, but rather forecast the inverse of the CAPE ratio itself. Specifically, we estimate a vector autoregressive (VAR) model with 12 monthly lags of the form

$$\boldsymbol{X}_t = \alpha + \beta_1 X_{t-1} + \beta_2 X_{t-2} + \ldots + \beta_{12} X_{t-12} + \varepsilon_t$$

where $\boldsymbol{X}_t$ is a vector of the five variables in the VAR model in logarithmic form, including

- CAPE real earnings yield, or 1/CAPE
- Real 10-year bond yields, or nominal Treasury yield less an estimated 10-year expected inflation rate (see the Appendix)
- Year-over-year CPI inflation rate
- Realized S&P 500 price volatility, over trailing 12 months
- Realized volatility of changes in our real bond yield series, over trailing 12 months.[3]

The motivation of including these five VAR variables derives from Asness [2003], who found that earnings yield rises when bond yields rise, stock volatility rises, and bond market volatility falls. Note that we lag the "E" in the CAPE ratio by six months and the CPI data by two months to account for real-time data availability at any month's end.

Step 2: Impute Stock Returns from the CAPE Earnings Yield Forecasts

Rather than estimating Equation (1), we calculate future returns directly based on their three components, thereby reducing estimation bias. We adapt the framework of Bogle and Nolan [2015] and Ferreira and Santa-Clara [2011] in imputing monthly stock returns by their sum of parts identity:

$$r_{t+1} \equiv \%\Delta PE_{t+1} + \%\Delta E_{t+1} + DP_{t+1} \qquad (2)$$

where %ΔPE is the percentage change in P/E ratio, %ΔE is earnings growth, and DP is the dividend yield. The VAR model's forecast for the earnings yield provides the percentage changes in CAPE ratios, $\%\Delta PE_{t+1}$, for imputing stock returns directly by the sum of parts (Equation (2)). For simplicity, we assume that earnings growth is constant and equal to its long-term average, whereas the dividend yield is the product of the earnings yield times the payout ratio.[4] As a result, only earnings yield (1/CAPE) has to be forecasted via regression to

Exhibit 6

Comparison of Different Stock-Forecasting Approaches

Stock Return Component	Ingredients in the Stock-Return Forecast: Traditional Shiller CAPE Ratio Regression	Bogle Occam's Razor Model	This Article's Two-Step Approach
Dividend yield	Swept into OLS alpha/ intercept coefficient	Actual dividend yield at beginning of period	Derived from forecasting earnings yield (below) times the beginning of period payout ratio
Earnings growth	Swept into OLS alpha/ intercept coefficient	Uses trailing 10-year earnings growth	Uses trailing long-run historical average earnings growth
CAPE ratio mean reversion process	Estimated by beta of the regression; not conditional on other variables	Linear proration over 10 years between last available P/E ratio and trailing 30-year average; provides speculative-return; not conditional on any other variables	Forecasted earnings yield from five-variable VAR model that also includes real bond yields, inflation, real bond volatility, and equity volatility

Source: Authors' calculations.

predict stock returns at a given horizon. At any point in time, the VAR forecasts the CAPE earnings yield out for 10 years and, via Step 2, derives an expected future 10-year-ahead return on U.S. stocks. Exhibit 6 summarizes the similarities and differences of our approach compared to (1) traditional Shiller CAPE regressions and (2) the Bogle Occam's razor model.

The potential benefit of our approach is that the fair-value CAPE ratio—to which the actual CAPE ratio should revert—is permitted to vary over time, conditional on the movements in these other fundamental variables.[5] It is this fair-value CAPE that should be the relevant benchmark for forecasting the equity risk premium, not the CAPE long-run average.[6] Put another way, if actual CAPE ratios revert back to our fair-value CAPE ratio and not to CAPE's historical average, then our two-step model should produce more reliable stock return forecasts than traditional Shiller CAPE regressions.

The VAR model dynamics for the earnings yield are intuitive. Exhibit 7 traces the impulse response function of the earnings yield (1/CAPE) to shocks to real bond yields. Movements in earnings yields are jointly determined by changes in real bond yields because both are measures of expected future economic growth and monetary policy.[7] The intuition for the positive correlation between real bond yields and stock earnings yield is simple: Lower expected economic growth implies lower real bond yields, which implies lower earnings yield on stocks, and thus a higher fair-value CAPE ratio, all else equal. The disinflationary period and bond-bull market of the 1980s coincided with rising stock valuations. As real interest rates fell below their historical averages in the 1990s and 2000s, equity earnings yield also remained below their own average levels.

Exhibit 7

Shocks to Real Bond Yields Lead to Higher CAPE Earnings Yields

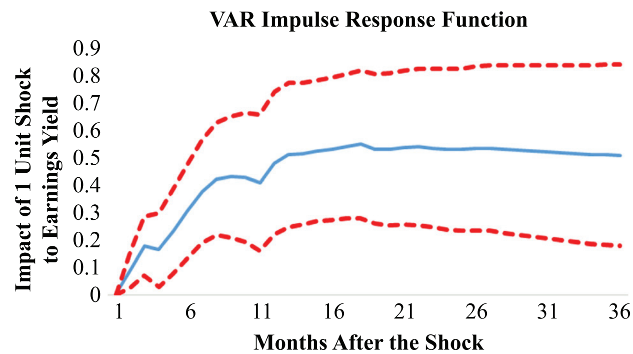

Note: Dotted lines reflect standard error bands.

Source: Authors' calculations.

COMPARING REAL-TIME FORECASTS: AN ILLUSTRATION

Exhibit 8 compares the projections for the earnings yield (1/CAPE) from two models: (1) that which is implied by a traditional Shiller CAPE regression[8] and (2) our VAR model. For convenience, we re-express the earnings yield as the CAPE ratio. We choose December 1999, when the CAPE was above 40, to illustrate relative forecast performance.

EXHIBIT 8
CAPE Real-Time Forecasts, VAR Model vs. Implied by Traditional Shiller Regressions

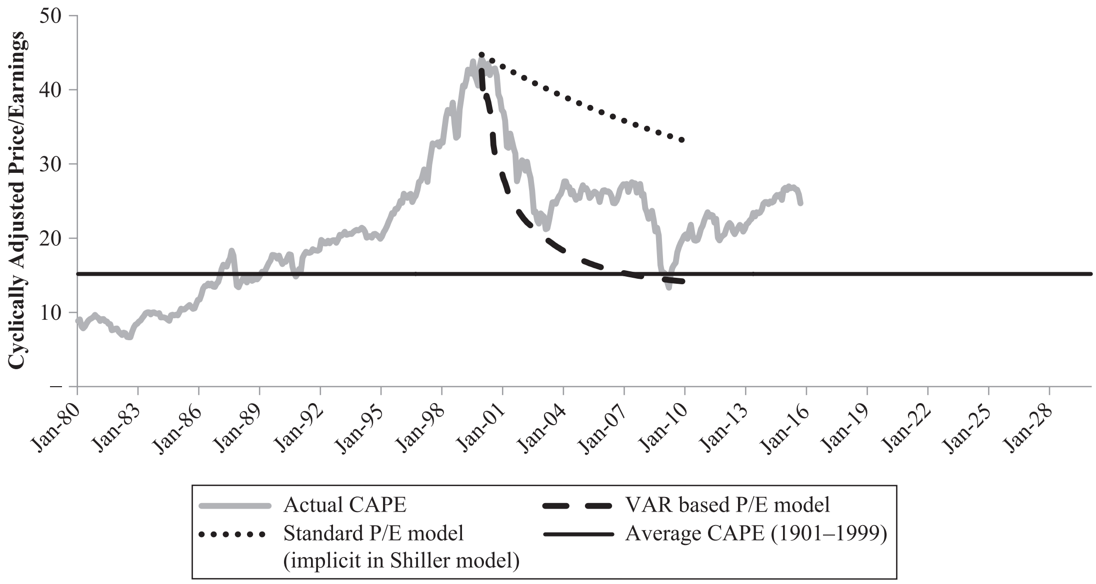

Note: Vanguard's model uses a VAR-based P/E model.
Source: Vanguard calculations, based on Robert Shiller's website, at aida.wss.yale.edu/~shiller/data.htm.

Following the dot-com bust period after 1999, the VAR-based CAPE projections anticipate subsequent CAPE trends more accurately than even Equation (1). This is because earnings yields are not assumed to converge unconditionally to their long-run average, as typical Shiller CAPE regressions do, but rather are a function of the current state of other variables in the model. Rising real bond yields—combined with the CAPE ratio above its fair value—leads to a sharper correction in earnings yields in our VAR model and hence to more accurate future stock market returns. This case study underscores why conditioning mean-reversion in CAPE ratios on real bond yields can improve long-run return forecasts.

COMPARING FORECAST PERFORMANCE

Exhibit 9 compares the predictability of our two-step approach to the Shiller and Siegel CAPE ratios by running out-of-sample forecasts for the U.S. stock market for the same two periods as before: 1960 to 2016 and 1985 to 2016. Like Philips and Ural [2016], we focus on nominal returns. Over both periods, our approach forecasts 10-year-ahead stock returns more accurately (in real time) for the United States than does the naïve constant-mean approach. More notable is that the two-step model's RMSE is lower and statistically different from the historical average for both Shiller and Siegel valuation metrics. Since 1985, the average forecast error (i.e., RMSE) using the two-step VAR model has been 4.1% compared to 7.8% for the Shiller CAPE ratio using Equation (1), a reduction of more than 40%.

Our results strongly suggest that conditioning mean reversion on macroeconomic conditions influences the steady-state CAPE ratio in real-time forecasts and thus leads to more accurate long-run return projections. This is demonstrated by the Diebold–Mariano tests of whether the VAR-based forecasts are better or worse than the historical mean. Since 1985, the errors from the VAR-based model's forecasts are statistically lower at the 1% significance level, whereas the other traditional models we mentioned are not. In fact, the forecasts from the traditional CAPE regression using the Shiller CAPE ratio are statistically worse at the 1% significance level. The VAR model performs best out of sample since 1960 using the Shiller CAPE ratio, but since 1985 using Siegel's CAPE ratio. We suggest in practical applications to use a simple weighted average of the two measures as a forecast diversification strategy.

Exhibit 9

Comparison of Real-Time Predictive Power, Nominal U.S. Stock Returns

	Out-of-Sample Forecasts Made Since 1960			Out-of-Sample Forecasts Made Since 1985		
Predictive Variable	**Correlation of Predicted Returns with Actual**	**Average Forecast Error (RMSE)**	**Model Forecast Error Relative to Error of Using a Naïve Historical Mean Forecast**	**Correlation of Predicted Returns with Actual**	**Average Forecast Error (RMSE)**	**Model Forecast Error Relative to Error of Using a Naïve Historical Mean Forecast**
Historical average		5.8%			6.2%	
Shiller CAPE ratio	83%	5.5%*	LOWER	91%	7.8%***	HIGHER
Siegel CAPE ratio	67%	4.9%***	LOWER	90%	6.2%	SIMILAR
Bogle Occam's razor	73%	4.7%***	LOWER	73%	5.9%	SIMILAR
Two-step VAR model, Shiller CAPE	81%	3.2%***	LOWER	90%	4.1%***	LOWER
Two-step VAR model, Siegel CAPE	67%	4.5%***	LOWER	90%	3.4%***	LOWER

Notes: The statistics shown are for 10-year annualized returns using the models described. An asterisk next to the RMSE refers to the significance of the Diebold–Mariano test [2002] of whether the forecast is statistically better or worse than the historical mean.

Significance levels of 90%, 95%, and 99% are denoted by one, two, and three stars, respectively.

Source: Authors' calculations based on the data sources listed in the Appendix.

Exhibit 10 shows the actual real-time forecast of our two-step model for U.S. stocks. Our fair-value CAPE approach tracks the actual rolling 10-year-ahead U.S. stock returns fairly well, declining throughout the 2000s and anticipating a strong rebound immediately following the global financial crisis in 2009. Traditional CAPE regressions are also highly correlated with future returns, yet they consistently project lower 10-year-ahead stock returns than what has been actually realized by investors over our sample period. Exhibit 11 shows that in contrast to the traditional CAPE models, our two-step approach exhibits better parameter stability in the out-of-sample period.

CONCLUSION

Valuation metrics such as P/E ratios are widely followed by the investment community because they are believed to predict future long-term stock returns. Arguably the most popular is Robert Shiller's CAPE, which is currently above its long-run average. However, the out-of-sample forecast accuracy of stock forecasts produced by CAPE ratios has become increasingly poor. In this article we have shown why and offer a solution to provide a more robust approach for producing long-run stock return forecasts.

Rather than focusing on the construction of the CAPE ratio, we propose here a new CAPE regression methodology that marries CAPE mean reversion with the current and expected future conditions in the macroeconomy. We find that a common industry approach of forecasting long-run stock returns can produce large forecast errors due to both estimation bias and its strict assumption that the CAPE ratio will revert over time to its long-run (and constant) mean. Although far from perfect, our model's out-of-sample forecasts for 10-year-ahead U.S. stock returns since 1960 are roughly 40%–50% more accurate than conventional methods. Real-time forecast differences in 10-year-ahead stock returns are statistically significant and have grown to exceed three percentage points after 1985, given the secular decline in real bond yields. In our model, lower real bond yields imply higher equilibrium CAPE ratios. This framework would appear to explain *both* elevated CAPE ratios and robust stock returns over the past two decades. Future research could involve testing our approach in non-U.S. markets with long-spanning data, or even sectors of the U.S. equity market.

Overall, we encourage investment professionals to explore (and extend) our VAR-based framework when forecasting stock returns for strategic asset allocation. As of June 2017, our model projects a guarded, lower-than-historical return on U.S. stocks of approximately 4.9% over the coming decade. This muted forecast for U.S. stock returns is not simply because the CAPE ratio is above its long-run mean.

EXHIBIT 10

Two-Step Fair-Value CAPE Model—Reasonable Out-of-Sample Performance

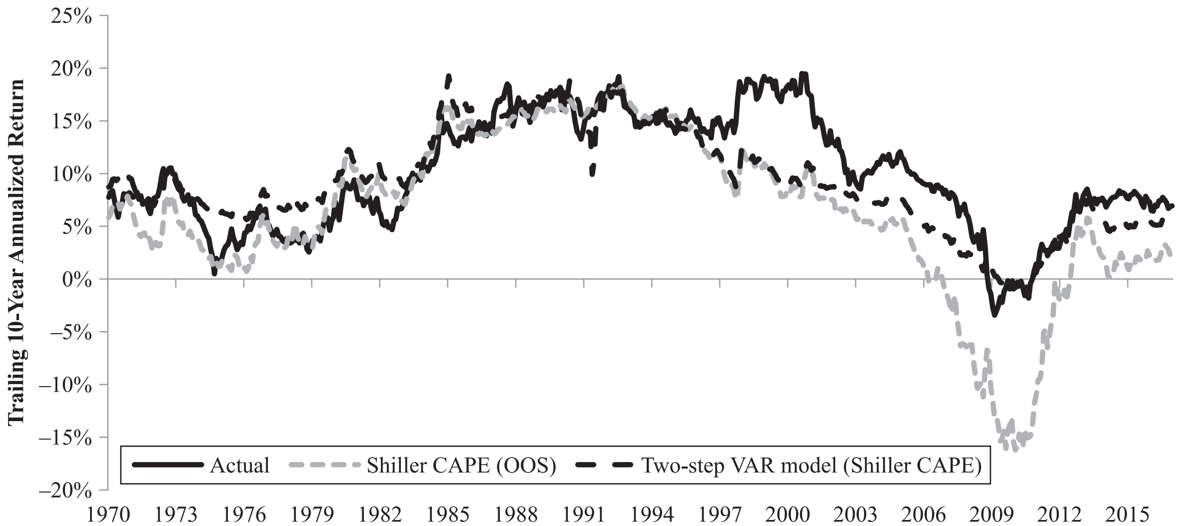

Note: For the real-time analysis, the regression coefficients are determined recursively at a monthly frequency, starting with January 1926–December 1959 data and re-estimating the regression coefficients every month thereafter.

Source: Authors' calculations.

EXHIBIT 11

Two-Step Fair-Value CAPE Model Appears More Stable

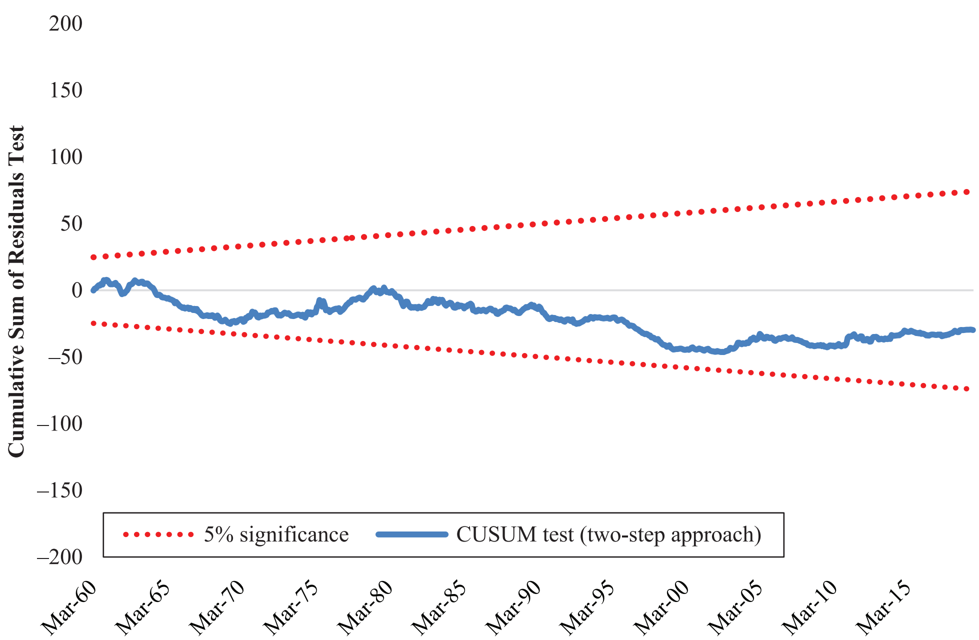

APPENDIX

EXHIBIT A1

Inflation Expectations and Real Yields

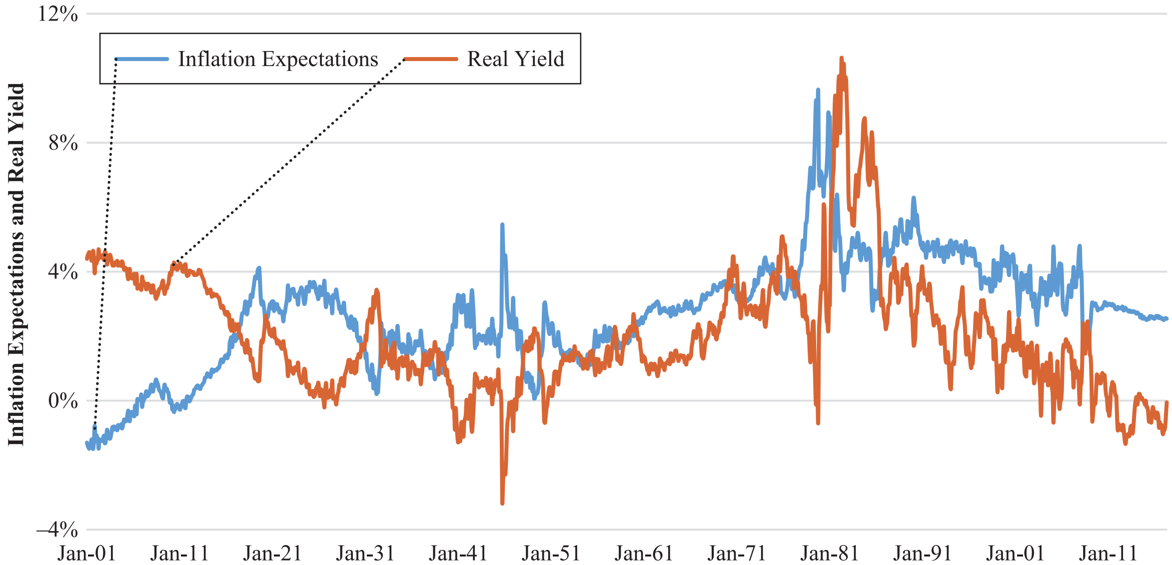

Notes: The model is an AR(12) model on monthly inflation with a 30-year rolling window. Initial estimation period is January 1871 through December 1900, after which monthly inflation is forecasted out for 10 years and annualized over 10 years to determine the inflation expectation in January 1901. The estimation window is rolled forward to estimate the inflation expectation series.

Source: Authors' calculations.

DATA APPENDIX

All of the data used in this article were obtained from Professor Robert Shiller's website, at aida.wss.yale .edu/~shiller/data.htm. Real bond yields reflect the nominal 10-year U.S. Treasury yield, less an estimate of 10-year-ahead inflation expectations. A consistently defined and long-running series on U.S. inflation expectations since the 1920s does not exist.

Our synthetic inflation expectations series was derived so that an investor could have replicated them at the time our stock forecasts were made. Specifically, we defined inflation expectations as the average of the predicted CPI inflation rate over the next 10 years generated from an AR model at any month in time. The AR model included 12 monthly lags in annualized CPI inflation rates and was estimated using a 30-year rolling window. The synthetic time series for our expected 10-year inflation rate is presented in Exhibit A1.

ENDNOTES

The views of the authors in this paper are not necessarily those of the Vanguard Group Inc. We wish to acknowledge Jack Bogle, Scott Pappas, Denis Chaves, and Fei Xu for their useful comments on this paper.

[1]The challenges of predicting stock returns over shorter horizons have been well documented by Campbell and Thompson [2008] and Goyal and Welch [2008]. For a survey of the literature on predicting the equity risk premium, see Ilmanen [2011], Damodaran [2012], and Davis, Aliaga-Diaz, and Thomas [2012].

[2]Inflation expectations appear less relevant and may explain why nominal bond yields (i.e., the Fed model) are not robust predictors of future stock returns. The results of our two-step model will illustrate this later.

[3]Arnott, Chaves, and Chow [2015] found that both real yields and inflation expectations are positively related to the earnings yield on U.S. stocks. It remains unclear why inflation expectations—a component of nominal bond yields—should influence earnings yields, considering stocks are a long-run inflation hedge. Importantly, this so-called *inflation illusion* effect is weaker in our VAR model than the effect from real bond yields given the joint dynamics of our VAR model, which we discuss later.

[4]The benefit of our sum of parts approach is that it should mitigate so-called Stambaugh [1999] bias that can

plague predictive regressions with persistent regressors like CAPE ratios that involve overlapping data (Nelson and Kim [1993]). In results that are unreported here but available upon request, including changes in earnings growth in the VAR does not materially alter the results. Consistent with the findings of Cochrane [2008], changes in earnings yields help predict future stock returns, not earnings growth.

[5]To predict future stock returns, we need not forecast the other five VAR factors accurately so much as account for their interdynamics in affecting earnings yields through time.

[6]The notion of a variable's unobserved fair value is common in macroeconomics. Examples include the full-employment concept of non-accelerating inflation rate of unemployment; a currency's purchasing power parity; and the natural rate of interest, or *R-star*.

[7]Historically, earnings yield and real bond yields have tended to move in tandem. We also know that breaks in real yields and inflation expectations occurred during the early 1950s (the end of the Treasury–Fed accord that pegged interest rates after World War II), the mid-1970s (the OPEC oil shock), and the early 1980s (when Volcker and the Fed tamped down higher inflation) given changes in macroeconomic regimes. In results unreported here, we link the structural breaks in Shiller CAPE regressions to breaks in real interest rates and other financial conditions that, when controlled for, should improve model stability.

[8]It can be shown that any predictive regression is equivalent to a single-period stock return equation plus an AR(1) or first-order regressive process in the predictor.

REFERENCES

Arnott, R.D., D.B. Chaves, and T. Chow. "King of the Mountain: Shiller P/E and Macroeconomic Conditions." SSRN, 2015. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2640327.

Asness, C.S. "Fight the Fed Model." *The Journal of Portfolio Management*, Vol. 30, No. 1 (2003), pp. 11-24.

Bogle, J.C. "Investing in the 1990s: Occam's Razor Revisited." *The Journal of Portfolio Management*, Vol. 18, No. 1 (1991), pp. 88-91.

——. "The 1990s at the Halfway Mark." *The Journal of Portfolio Management*, Vol. 21, No. 4 (1995), pp. 21-31.

Bogle, J.C., and M.W. Nolan, Jr. "Occam's Razor Redux: Establishing Reasonable Expectations for Financial Market Returns." *The Journal of Portfolio Management*, Vol. 42, No. 1 (2015), pp. 119-134.

Brown, R.L., J. Durbin, and J.M. Evans. "Techniques for Testing the Constancy of Regression Relationships over Time." *Journal of the Royal Statistical Society, Series B (Methodological)*, Vol. 37, No. 2 (1975), pp. 149-192.

Campbell, J.Y., and R.J. Shiller. "Stock Prices, Earnings, and Expected Dividends." *The Journal of Finance*, Vol. 43, No. 3 (1988), pp. 661-676.

——. "Valuation Ratios and the Long-Run Stock Market Outlook." *The Journal of Portfolio Management*, Vol. 24, No. 2 (1998), pp. 11-26.

Campbell, J.Y., and S.B. Thompson. "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?" *The Review of Financial Studies*, Vol. 21, No. 4 (2008), pp. 1509-1531.

Cochrane, J.H. "The Dog that Did Not Bark: A Defense of Return Predictability." *The Review of Financial Studies*, Vol. 21, No. 4 (2008), pp. 1533-1575.

Damodaran, A. "Equity Risk Premiums (ERP): Determinants, Estimation and Implications." SSRN, 2012. http://ssrn.com/abstract=2027211.

Davis, J.H., R. Aliaga-Diaz, and C.J. Thomas. "Forecasting Stock Returns: What Signals Matter, and What Do They Say Now?" Investment Strategy Group White Paper, Vanguard, 2012.

Diebold, F.X., and R.S. Mariano. "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics*, Vol. 20, No. 1 (2002), pp. 134-144.

Ferreira, M.A., and P. Santa-Clara. "Forecasting Stock Market Returns: The Sum of the Parts Is More than the Whole." *Journal of Financial Economics*, Vol. 100, No. 3 (2011), pp. 514-537.

Goyal, A., and I. Welch. "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction." *The Review of Financial Studies*, Vol. 21, No. 4 (2008), pp. 1455-1508.

Ilmanen, A. *Expected Returns: An Investor's Guide to Harvesting Market Rewards.* West Sussex, United Kingdom: John Wiley & Sons, 2011.

Nelson, C.R., and M.J. Kim. "Predictable Stock Returns: The Role of Small Sample Bias." *The Journal of Finance*, Vol. 48, No. 2 (1993), pp. 641-661.

Pettenuzzo, D. and A.G. Timmermann. "Predictability of Stock Returns and Asset Allocation Under Structural Breaks." *Journal of Econometrics*, Vol. 164, No. 1 (2011), pp. 60-78.

Philips, T., and C. Ural. "Uncloaking Campbell and Shiller's CAPE: A Comprehensive Guide to Its Construction and Use." *The Journal of Portfolio Management*, Vol. 43, No. 1 (2016), pp. 109-125.

Siegel, J.J. "The Shiller CAPE Ratio: A New Look." *Financial Analysts Journal*, Vol. 72, No. 3 (2016), pp. 41-50.

Stambaugh, R.F. "Predictive Regressions." *Journal of Financial Econometrics*, Vol. 54, No. 3 (1999), pp. 375-421.