the journal of
PORTFOLIO
management

The Stock-Bond Correlation

Megan Czasonis, Mark Kritzman, and David Turkington

Megan Czasonis

Megan Czasonis is a Managing Director for the Portfolio and Risk Research team at State Street Associates. The Portfolio and Risk Research team collaborates with academic partners to develop new research on asset allocation, risk management, and investment strategy. The team delivers this research to institutional investors through indicators, advisory projects, and thought leadership pieces. Megan has co-authored various journal articles and works closely with institutional investors to develop customized solutions based on this research. Megan graduated Summa Cum Laude from Bentley University with a B.S. in Economics / Finance.

Mark Kritzman

Mark Kritzman is a Founding Partner and CEO of Windham Capital Management, a Founding Partner of State Street Associates, and teaches a graduate finance course at the Massachusetts Institute of Technology. Mark served as a Founding Director of the International Securities Exchange and has served on several Boards, including the Institute for Quantitative Research in Finance, The Investment Fund for Foundations, and State Street Associates. He has written numerous articles for academic and professional journals and is the author / co-author of seven books including *A Practitioner's Guide to Asset Allocation*, *Puzzles of Finance*, and *The Portable Financial Analyst*. Mark won Graham and Dodd Scrolls in 1993 and 2002, the Research Prize from the Institute for Quantitative Investment Research in 1997, the Bernstein Fabozzi/Jacobs Levy Award nine times, the Roger F. Murray Prize from the Q-Group in 2012, and the Peter L. Bernstein Award in 2013 for Best Paper in an Institutional Investor Journal. Mark has a BS in economics from St. John's University, an MBA with distinction from New York University, and a CFA designation.

David Turkington

David Turkington is Senior Managing Director and Head of Portfolio and Risk Research at State Street Associates. His team is responsible for research and advisory spanning asset allocation, risk management and quantitative investment strategy. Mr. Turkington is a frequent presenter at industry conferences, has published research articles in a range of journals, and led the development of State Street's systemic risk and turbulence indicators. He is also co-author of the book A Practitioner's Guide to Asset Allocation. His research has received the 2013 Peter L. Bernstein Award, four Bernstein-Fabozzi/Jacobs-Levy Outstanding Article Awards, and the 2010 Graham and Dodd Scroll Award. Mr. Turkington graduated summa cum laude from Tufts University with a BA in mathematics and quantitative economics, and he holds the CFA designation.

The Stock-Bond Correlation

Megan Czasonis, Mark Kritzman, and David Turkington

Megan Czasonis
is a managing director at State Street Associates in Cambridge, MA.
mczasonis@statestreet.com

Mark Kritzman
is the chief executive officer at Windham Capital Management in Boston, MA, and a senior lecturer at the MIT Sloan School of Management in Cambridge, MA.
kritzman@mit.edu

David Turkington
is a senior managing director at State Street Associates in Cambridge, MA.
dturkington@statestreet.com

KEY FINDINGS

- The stock-bond correlation is a critical component of many investment activities, such as forming optimal portfolios, designing hedging strategies, and assessing risk.

- Most investors estimate the correlation of longer-interval returns by extrapolating the correlation of past shorter-interval returns, but this approach is decidedly unreliable.

- By applying recent advances in quantitative methods, it is possible to generate reliable predictions of the correlation of longer-horizon stock and bond returns.

ABSTRACT

Investors rely on the stock-bond correlation for a variety of tasks, such as forming optimal portfolios, designing hedging strategies, and assessing risk. Most investors estimate the stock–bond correlation simply by extrapolating the historical correlation of monthly returns; they assume that this correlation best characterizes the correlation of future annual or multiyear returns, but this approach is decidedly unreliable. The authors introduce four innovations for generating a reliable prediction of the stock-bond correlation. First, they show how to represent the correlation of single-period cumulative stock and bond returns in a way that captures how the returns drift during the period. Second, they identify fundamental predictors of the stock-bond correlation. Third, they model the stock–bond correlation as a function of the path of some fundamental predictors rather than single observations. Finally, they censor their sample to include only relevant observations, in which relevance has a precise mathematical definition.

TOPICS

*Portfolio management/multi-asset allocation, risk management, statistical methods**

Investors rely on the stock-bond correlation for a variety of activities, such as forming optimal portfolios, designing hedging strategies, and assessing risk. For these activities, investors are usually, or at least often, interested in the correlation of longer-horizon returns such as annual or multiyear returns. Most investors predict the correlation of longer-interval stock and bond returns by extrapolating the historical correlation of shorter-interval returns, but this approach is decidedly unreliable. We introduce several innovations for generating a reliable prediction of the correlation of longer-interval stock and bond returns, and we document each innovation's marginal contribution to improving the predictive power of our approach.

We first introduce the notion of a single-period correlation. This innovation remedies the dual problem that correlations estimated from shorter-interval returns differ significantly from correlations estimated from longer-interval returns for the same

EXHIBIT 1
Stock–Bond Correlation (January 1926–November 2019)
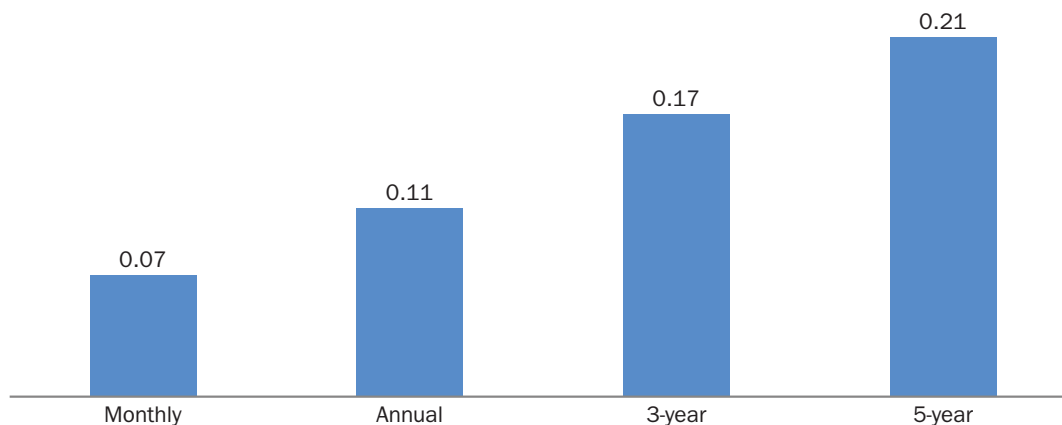
NOTE: The correlations of annual, 3-year, and 5-year returns are estimated from overlapping monthly observations.

historical sample and that the correlation of longer-interval returns varies through time. Next, we identify fundamental predictors of the stock-bond correlation, which capture conditions both in the real economy and the financial markets. We then address the fact that the stock-bond correlation depends not only on point-in-time observations but also on the path of certain economic variables as well. Finally, we apply a recent innovation called partial sample regression, which isolates relevant observations from the historical sample and uses a crafty mathematical equivalence to derive a better prediction of the stock-bond correlation.

Although this overall process is complex, we demonstrate its value by documenting the incremental improvement in forecast reliability as we proceed from naïve extrapolation (which is common practice) to the full application of all these innovations.

SINGLE-PERIOD CORRELATION

Our focus is to forecast the correlation of stock and bond returns for annual or multiyear horizons; we assume that investors care less about how the main growth and defensive components of their strategies co-move month to month and more about how they co-move over the duration of the investment horizon. To illustrate our approach, we focus on the correlation of five-year returns. We use the S&P 500 Index to estimate stock returns and the Bloomberg US Long Treasury Index to estimate bond returns.[1]

The conventional approach for forecasting the correlation of five-year stock and bond returns is to extrapolate the correlation of monthly returns from a recent historical sample. This approach assumes that correlations are invariant to the return interval used to estimate them. However, this invariance would only be true if the returns of stocks and bonds were each serially independent at all lags and if their lagged cross correlations were also zero at all lags. These implicit assumptions are not borne out by historical evidence, as shown in Exhibit 1.

Exhibit 1 offers stark evidence that the stock-bond correlation is far from stable across return intervals. This instability may arise, for example, if short-term returns

---

[1] Prior to 1973, we use the Ibbotson Stock Index and the Ibbotson Treasury Bond Index to estimate stock and bond returns, respectively.

respond similarly to a given factor, whereas longer-term returns drift apart in response to a different, lower frequency influence. Equation 1 shows how the correlation of longer-interval returns is mathematically related to the correlation of shorter-interval returns.

$$\rho(x_t + \cdots + x_{t+q-1},\ y_t + \cdots + y_{t+q-1}) = \frac{q\rho_{x_t,y_t} + \sum_{k=1}^{q-1}(q-k)(\rho_{x_{t+k},y_t} + \rho_{x_t,y_{t+k}})}{\sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\rho_{x_t,x_{t+k}}}\ \sqrt{q + 2\sum_{k=1}^{q-1}(q-k)\rho_{y_t,y_{t+k}}}}$$

(1)

where $\rho(x_t + \cdots + x_{t+q-1},\ y_t + \cdots + y_{t+q-1})$ is the correlation of the cumulative continuous returns of *x* and *y* over *q* periods.

The numerator of Equation 1 equals the covariance of the assets, taking lagged cross correlations into account; the denominator equals the product of the assets' standard deviations, taking their autocorrelations into account. To apply this equation, we would need to estimate the autocorrelations of stocks and bonds at all lags as well as their lagged cross correlations at all lags, which would be unduly cumbersome.[2]

Alternatively, we can calculate the correlation of independent five-year returns, but this approach has two drawbacks. First, the estimate of the correlation is highly sensitive to the start date of the first five-year observation. And second, the stock-bond correlation is unlikely to be constant across such a long history. These problems might lead us to consider using overlapping observations to calculate the stock-bond correlation, but this approach, although mitigating the start-date problem, is still subject to issues of time variation.

Therefore, the first order of business is to develop an approximation of the correlation of five-year stock and bond returns that captures the extent to which these returns move synchronously or drift apart during each five-year period. Just as each five-year return for stocks or bonds differs from the long-run average returns, the co-movement of stocks and bonds during each five-year period differs from the co-movement on average over the long run.

Equation 2 captures the co-movement of the cumulative return of stocks and bonds for a chosen five-year period, considering the average co-movement of stocks and bonds over the full sample of five-year returns.

$$Single\text{-}period\ correlation_t(s,b) = \frac{\left(\dfrac{r_{s,t} - \mu_s}{\sigma_s}\right)\left(\dfrac{r_{b,t} - \mu_b}{\sigma_b}\right)}{\dfrac{1}{2}\left(\left(\dfrac{r_{s,t} - \mu_s}{\sigma_s}\right)^2 + \left(\dfrac{r_{b,t} - \mu_b}{\sigma_b}\right)^2\right)}$$
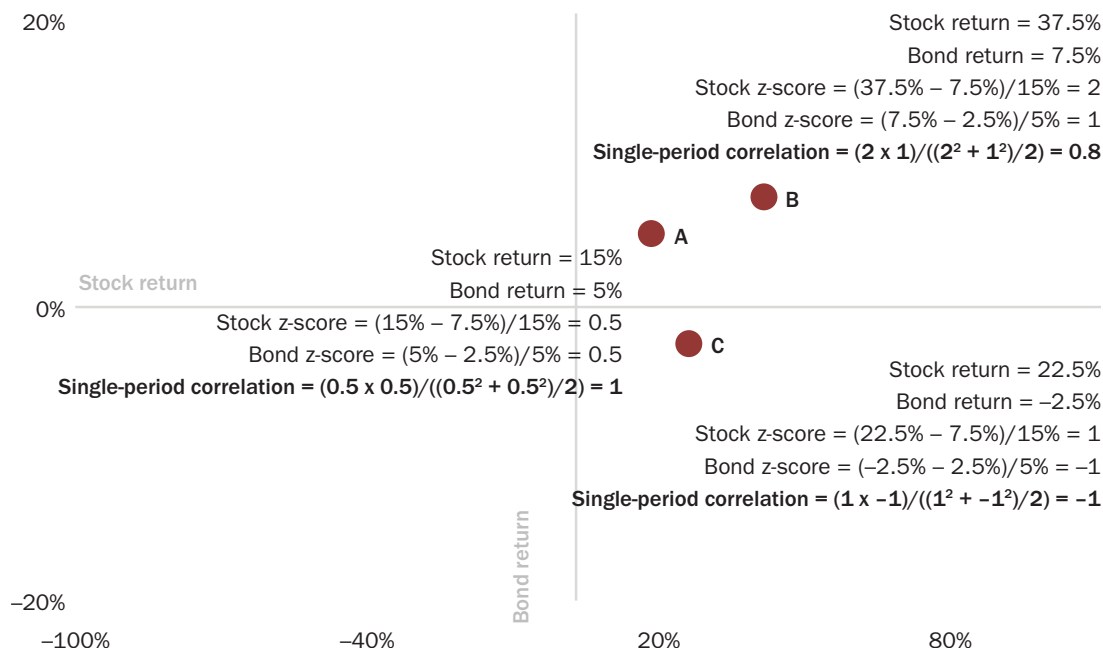
(2)

In Equation 2, r_s and r_b equal the five-year cumulative return of stocks and bonds, respectively; μ_s and μ_b equal the long-run arithmetic average return of stocks and bonds, respectively; and σ_s and σ_b equal the standard deviation of the five-year stock and bond returns, respectively.

We refer to this measure as a single-period correlation because it measures the co-movement of the cumulative returns of stocks and bonds for a single five-year period rather than the average correlation of five-year returns over the full sample. Each single-period correlation will differ from the long-run average correlation based on the period-specific co-movement of stocks and bonds.

---

[2] See Kinlaw, Kritzman, and Turkington (2014) for more detail about the divergence of high- and low-frequency estimation of correlations.

opposed to noise. Later, we will refer to the concept of informativeness again, but in the broader context of linear regression with independent variables.[3]

If, in Equation 3, we rewrite informativeness as the average squared z-score of stock and bond returns and substitute Equation 2 for single-period correlation, the correlation then becomes

$$Correlation\,(s,b) = \frac{1}{N}\sum_{t=1}^{N}\left[\frac{1}{2}\left(\left(\frac{r_{s,t}-\mu_s}{\sigma_s}\right)^2+\left(\frac{r_{b,t}-\mu_b}{\sigma_b}\right)^2\right)\times\frac{\left(\frac{r_{s,t}-\mu_s}{\sigma_s}\right)\left(\frac{r_{b,t}-\mu_b}{\sigma_b}\right)}{\frac{1}{2}\left(\left(\frac{r_{s,t}-\mu_s}{\sigma_s}\right)^2+\left(\frac{r_{b,t}-\mu_b}{\sigma_b}\right)^2\right)}\right]$$

(4)

This expression simplifies to the average product of stock and bond z-scores over the full sample, which is equivalent to the full-sample covariance of stock and bond returns divided by the product of their full-sample standard deviations, the formula commonly used to measure a time series correlation.

$$Correlation\,(s,b) = \frac{1}{N}\sum_{t=1}^{N}\left(\left(\frac{r_{s,t}-\mu_s}{\sigma_s}\right)\left(\frac{r_{b,t}-\mu_b}{\sigma_b}\right)\right) = \frac{Cov(s,b)}{\sigma_s\sigma_b}$$

(5)

At this point, we believe we have made our case that the single-period correlation is a better representation of the correlation of five-year returns than the correlation of monthly returns; it captures the drift of stock and bond returns, whereas the correlation of monthly returns does not. It is also better than the correlation of either independent blocks of five-year returns or overlapping five-year returns because it captures the time-varying nature of the stock-bond correlation whereas the full-sample correlation does not.

We now move on to the fundamental predictors of the stock-bond correlation.

FUNDAMENTAL PREDICTORS OF THE STOCK–BOND CORRELATION

We suggest that the correlation of longer-term stock and bond returns responds to fundamental factors that influence the gradual drift in stock and bond returns.[4] We have identified four fundamental factors to predict the correlation of five-year stock and bond returns. We select economic growth and inflation because they are key drivers of stock and bond returns, respectively, and the interplay between them may affect the stock-bond correlation. For example, a positive demand shock will increase both growth and inflation, which is likely to induce a divergence in stock and bond performance. By contrast, a supply-side shock, such as the 1970s oil crisis, may lead to higher inflation but lower growth, which is likely to push stock and bond prices in the same direction. We also select the relative level and volatility of stock and bond yields because they likely affect investor demand for stocks versus bonds. For example, when stock and bond yields are similar, both assets are similarly attractive, implying

---

[3] In linear regression, informativeness takes into account the correlations of the independent variables with each other because these effects are distinct from the dependent variable, which is the object of the analysis. In the context of estimating the Pearson correlation, informativeness only takes into account the magnitude of individual variables because the correlation itself is the object of the analysis.

[4] Previous research has also suggested using fundamental factors to generate a forward-looking view of the stock-bond correlation. For example, see Johnson et al. (2014).

a positive stock-bond correlation. However, when stock yields are meaningfully higher or riskier than bond yields, we might expect a negative stock-bond correlation.

We define each factor for the US market as described in the following. As we mentioned earlier, we use the S&P 500 Stock Index to represent stocks and the Bloomberg US Long Treasury Bond Index to represent bonds.

Fundamental Factors

- Economic growth (year-over-year percentage change in industrial production)
- Inflation (year-over-year percentage change in the Consumer Price Index)
- Relative yield of stocks and bonds (natural log of the earnings yield divided by the 10-year Treasury rate)[5]
- Relative volatility of stock and bond yield (natural log of the five-year trailing volatility of the earnings yield divided by the five-year trailing volatility of the 10-year Treasury rate)

Model Specification

We define our variable of interest as the subsequent five-year single-period correlation because, as we discussed, we believe that investors care more about how stocks and bonds co-move throughout the course of the investment horizon than how they co-move from month to month. We test five models to predict the subsequent five-year single-period correlation. In all cases, we use data starting in January 1926 to forecast the single-period five-year correlation beginning in January 1930 through December 2018, and we update the predictions annually.

Prior to running the regression-based models (Models 2 through 5, below), we transform the dependent variable to prevent the models from predicting a correlation that violates the bounds of negative and positive one. First, we shift and re-scale the single-period correlation so that it ranges from zero to one:

$$SPC^*(s,b) = \frac{1 + \textit{single-period correlation}\,(s,b)}{2} \tag{6}$$

Here, $SPC^*(s,b)$ is the shifted and re-scaled single-period correlation ranging from zero to one. Then, we apply the probit transformation so that it is unbound and normally distributed:

$$y = probit\,(SPC^*(s,b)) = \Phi^{-1}(SPC^*(s,b)) \tag{7}$$

In Equation 7, y is the probit of the shifted and re-scaled correlation, $SPC^*(s,b)$, estimated as the inverse of the cumulative distribution function of the standard normal distribution with probability $SPC^*(s,b)$.

After running the regressions, we transform the resulting predictions back so that they are correlations ranging from negative to positive one:

$$\hat{Y} = 2\Phi(\hat{y}) - 1 \tag{8}$$

In this equation, $\hat{Y}$ is the predicted single-period correlation ranging from negative to positive one, $\hat{y}$ is the predicted probit of the shifted and re-scaled correlation, and Φ denotes the cumulative distribution function of the standard normal distribution.

---

[5]We estimate the earnings yield as the inverse of the cyclically adjusted price/earnings ratio obtained from Robert J. Shiller's website and the 10-year Treasury rate as the 10-year Treasury constant maturity rate obtained from the Federal Reserve Bank of St. Louis (prior to 1953, we use the 10-year Treasury rate from Robert J. Shiller's website).

Model 1: naïve extrapolation of trailing correlation. This model simply extrapolates the correlation of monthly stock and bond returns from the prior five-year period to predict the single-period correlation for the next five-year period. Not only is this approach the simplest way to model the stock-bond correlation, it is by far the most commonly used approach.

Model 2: regression on trailing correlation. In this model, we regress the transformed five-year single-period correlation on the correlation of monthly stock and bond returns from the prior five-year period. This approach differs from naïve extrapolation in that the intercept and slope will differ from 0 and 1, respectively, which is implied by naïve extrapolation.

Model 3: regression on fundamental factors. In this model, we regress the transformed five-year single-period correlation on four fundamental factors: industrial production, inflation, relative yield of stocks and bonds, and relative volatility of stock and bond yields, as defined earlier. The predictors are measured as of the end of the period prior to the estimation of the single-period correlation.

Model 4: regression on fundamental factors filtered for relevance. In this model, we use the same predictors we used in Model 3, but we filter the historical observations for their relevance, based on a technique called partial sample regression.[6]

Partial sample regression relies on a mathematical equivalence. The prediction from linear regression is a function of the weighted average of the past values of the dependent variable in which the weights are the relevance of the past observations for the independent variables. *Relevance* within this context has a precise meaning. It is the sum of the statistical similarity of the past observations to the current values for the independent variables, which is the negative of their Mahalanobis distances, and the informativeness of the past observations, which equals their Mahalanobis distances from the average values of the independent variables.

Equation 9 measures the multivariate similarity between x_i and x_t. It is simply the opposite (negative) of the multivariate distance between x_i and x_t.

$$Similarity\,(x_i, x_t) = -(x_i - x_t)\Omega^{-1}(x_i - x_t)' \tag{9}$$

Here x_t is a vector of the current values of the independent variables, x_i is a vector of the prior values of the independent variables, the symbol $'$ indicates matrix transpose, and Ω^{-1} is the inverse covariance matrix of X where X comprises all the vectors of the independent variables. This measure takes into account not only how independently similar the components of the x_is are to those of the x_ts, but also the similarity of their co-occurrence to the co-occurrence of the x_ts. All else equal, prior observations for the independent variables that are more like the current observations are more relevant than prior observations that are less similar.

However, not all measurements of similarity are alike. Observations that are close to their historical averages may be driven more by noise than by events. These ordinary occurrences are therefore less relevant. Observations that are distant from their historical averages are unusual and therefore more likely to be driven by events. These event-driven observations are potentially more informative.[7] Given this intuition, we define the informativeness of a prior observation x_t as its multivariate distance from its average value, $\bar{x}$. This definition of informativeness is the same as used previously to compute the Pearson correlation, but we now include the off-diagonal elements of the covariance matrix because they pertain to X variables, which are independent from our object of focus, y.

---

[6] We base our predictions on a subset of the 25% most relevant observations.

[7] For further discussion of noise-driven versus event-driven observations and their relationship to estimating risk, see Chow et al. (1999).

EXHIBIT 3
Correlation of Model Predictions and Actual Correlations
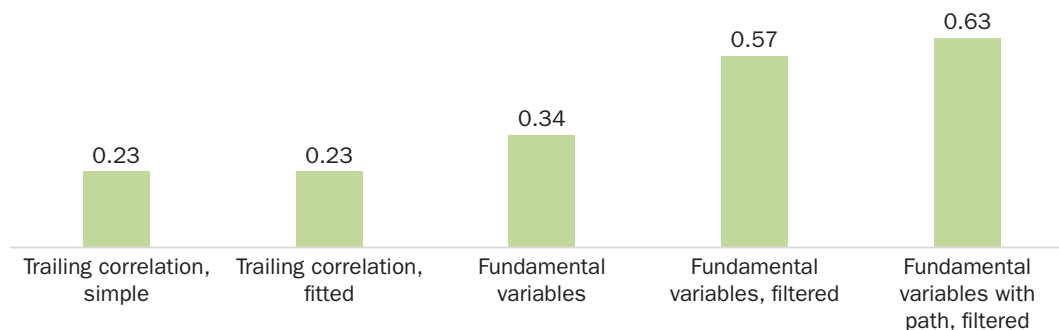
NOTES: By construction, models with more predictors are likely to have greater explanatory power, even if those predictors reflect random noise. The adjusted R^2 is often used to address this issue; however, we cannot apply it to partial sample regression owing to complexity in estimating the model's degrees of freedom. Therefore, we use simulation to evaluate the explanatory power of each model beyond what is expected from purely random effects. Specifically, we shift the time series of Y so that it is aligned in arbitrary ways with the predictors. This maintains the empirical distribution and serial dependence of the variables. We then fit each model and record the resulting correlation of predictions with simulated (shifted) Y_s. We repeat this process 82 times (the total number of possible offsets in Y) to create a distribution of random correlations for each model. We evaluate the quality of a model's fit according to where its actual correlation falls within its distribution of random correlations. The results suggest meaningful improvement from filtering by relevance. In the case of the *Fundamental variables, filtered* model, only five random correlations exceeded the actual result (corresponding to a p-value of 0.06). While the filtered model with paths and full sample fundamental model have greater explanatory power than trailing correlation, their results are not as statistically significant (*Trailing Correlation* p-values = 0.27, *Fundamental variables* p-value = 0.44, *Fundamental variables with path, filtered* p-value = 0.37). It is useful to evaluate these results alongside the Henriksson-Merton results (Exhibit 4), which are less likely to be dominated by noise of modest differences around correlations of zero.

$$Informativeness(x_i) = (x_i - \overline{x})\Omega^{-1}(x_i - \overline{x})' \tag{10}$$

The relevance of an observation x_i is equal to the sum of its multivariate similarity and its informativeness.

$$Relevance(x_i) = Similarity(x_i, x_t) + Informativeness(x_i) \tag{11}$$

To summarize, similarity equals the negative of the Mahalanobis distance of a prior observation of x_i from its current observation x_t. Informativeness equals the Mahalanobis distance of x_i from its historical average. Relevance equals the sum of similarity and informativeness. In other words, prior periods that are like the current period but are different from the historical average are more relevant than those that are not.

Once we determine the relevance of the past observations of the independent variables, we rely on the equivalence we noted earlier—that the prediction from linear regression is a function of the weighted average of the past values of the dependent variable in which the weights are the relevance of the past observations for the independent variables. This equivalence allows us to use Equations 12 and 13 to derive a new prediction for the stock-bond correlation from a subset of relevant observations.[8]

$$\hat{y}_t = \overline{y} + \frac{1}{2n}\sum_{i=1}^{n}(Similarity(x_i, x_t) + Informativeness(x_i))(y_i - \overline{y}) \tag{12}$$

---

[8] See Czasonis, Kritzman, and Turkington (2020) for a thorough discussion of this technique.

EXHIBIT 4
Henriksson-Merton Scores, Positive versus Negative

	Level		Change	
Prediction Models	**Statistic**	***P*-value**	**Statistic**	***P*-value**
Trailing correlation of monthly returns, extrapolated	1.06	0.29	1.13	0.19
Trailing correlation of monthly returns, fitted	1.20	0.06	1.13	0.19
Fundamental variables	1.25	0.04	1.18	0.12
Fundamental variables, filtered for relevance	1.27	0.01	1.18	0.12
Fundamental variables, paths and filtered for relevance	1.29	0.00	1.41	0.01

$$\bar{y} = \frac{1}{n}\sum_{i=1}^{n} y_i \qquad (13)$$

Model 5: regression on the path of fundamental factors filtered for relevance. In this model we redefine industrial production and inflation as paths. Instead of including just the prior periods' values for industrial production and inflation, we include values for the prior five periods, along with the prior period values for the relative stock and bond yield and the relative volatility of the stock and bond yield. Therefore, the five-year single-period correlation in period $t + 1$ is regressed on industrial production and inflation for periods t, $t - 1$, $t - 2$, $t - 3$, and $t - 4$, and relative stock and bond yield and relative volatility of the stock and bond yield in period t. Moreover, we filter the historical observations for their relevance, as we did for Model 4.

Exhibit 3 shows the correlation of the model predictions and the actual subsequent single-period correlations of five-year returns.

Exhibit 3 reveals some interesting insights. First, the most common approach for forecasting the stock-bond correlation, which is to extrapolate the correlation of monthly returns, ties as the least reliable approach. It is tied with regressing the five-year single-period correlation on the correlation of monthly stock and bond returns from the prior five-year period. Fitting the observations shifts the intercept and slope away from 0 and 1, respectively, but it does not improve the explanatory power. The introduction of fundamental variables to predict the stock-bond correlation significantly improves explanatory power, but the largest improvement comes from filtering the historical observations for their relevance. There is further benefit to replacing single observations for industrial production and inflation with paths.

Next, we evaluate these models by applying a nonparametric procedure, called the Henriksson-Merton test, to determine how well these models anticipate whether the subsequent level of the stock-bond correlation will be positive or negative.[9] Specifically, we calculate the percentage of times each model predicted the subsequent correlation would be positive when it was positive and that it would be negative when it was negative. We also test to determine how well these models predict whether the change in the subsequent correlation will be positive or negative. These percentages would sum to 2 if a model perfectly predicted the sign of the stock-bond correlation level or change. If a model had no ability to determine whether the subsequent level or change in correlation would be positive or negative, these percentages would sum to 1. Any value between 1 and 2 would be indicative of some degree of predictive power. A value below 1 would indicate that a model was worse than random or, more forgivingly, a reverse indicator. Exhibit 4 shows the Henriksson–Merton scores for the five models we tested.

---

[9] This nonparametric procedure was introduced by Henriksson and Merton (1981).

Exhibit 4 reveals that the innovations described earlier are, on balance, beneficial to forecasting the subsequent sign of both the level and change in the stock-bond correlation. In combination, they provide the most reliable forecasts.

SUMMARY

Investors rely on the stock-bond correlation to construct optimal portfolios, design hedging strategies, and assess risk. For many, if not most, applications investors care less about how stocks and bonds co-move month to month than about their co-movement over the duration of the investment horizon. The most common approach for estimating the longer-term correlation of stocks and bonds is to extrapolate the correlation of monthly returns over a prior period. However, this approach is decidedly unreliable. We introduce several innovations for forecasting the longer-term correlation of stocks and bonds. We introduce the notion of a single-period correlation to address the problem of the autocorrelations and lagged cross-correlations of stock and bond returns being nonzero and that longer-horizon correlations vary through time. We then identify several fundamental variables to predict the longer-horizon stock-bond correlation. We include paths as well as single observations in our regressions, and we filter historical observations for their statistical relevance. We find that these innovations significantly improve the reliability of our forecast of the stock-bond correlation.

Although we have only addressed the correlation of stock and bond returns in this article, we have no reason to doubt that the methodology proposed herein would work as well with other asset classes. It may be the case that the appropriate fundamental predictors would vary with the asset classes under consideration. Moreover, the degree to which our approach can be successfully transferred to other asset class pairs is an empirical issue. Nonetheless, the conceptual basis for our approach should apply broadly.

REFERENCES

Chow, G., E. Jacquier, M. Kritzman, and K. Lowry. 1999. "Optimal Portfolios in Good Times and Bad." *Financial Analysts Journal* 55 (3) (May/June): 65–71.

Czasonis, M., M. Kritzman, and D. Turkington. 2020. "Addition by Subtraction: A Better Way to Predict Factor Returns—And Everything Else." *The Journal of Portfolio Management* 46 (8): 98–107.

Henriksson, R., and R. Merton. 1981. "On Market Timing and Investment Performance. II. Statistical Procedures for Evaluating Forecasting Skills." *The Journal of Business* 54 (4): 513–533.

Johnson N., V. Naik, S. Page, N. Pederson, and S. Sapra. 2014. "The Stock-Bond Correlation." *Journal of Investment Strategies* 4 (1) (December): 3–18.

Kinlaw, W., M. Kritzman, and D. Turkington. 2014. "The Divergence of High- and Low-Frequency Estimation: Causes and Consequences." *The Journal of Portfolio Management* 40 (50): 156–168.

To order reprints of this article, please contact David Rowe at d.rowe@pageantmedia.com *or 646-891-2157.*