

May 08, 2013

832

Are Stocks Cheap? A Review of the Evidence

Fernando Duarte and Carlo Rosa


We surveyed banks, we combed the academic literature, we asked economists at central banks. It turns out that most of their models predict that we will enjoy historically high excess returns for the S&P 500 for the next five years. But how do they reach this conclusion? Why is it that the equity premium is so high? And more importantly: Can we trust their models?

The equity risk premium is the *expected* future return of stocks minus the risk-free rate over some investment horizon. Because we don't directly observe market expectations of future returns, we need a way to figure them out indirectly. That's where the models come in. In this post, we analyze twenty-nine of the most popular and widely used models to compute the equity risk premium over the last fifty years. They include surveys, dividend-discount models, cross-sectional regressions, and time-series regressions, which together use more than thirty different variables as predictors, ranging from price-dividend ratios to inflation. Our calculations rely on real-time information to avoid any look-ahead bias. So, to compute the equity risk premium in, say, January 1970, we only use data that was available in December 1969.

Let's now take a look at the facts. The chart below shows the weighted average of the twenty-nine models for the one-month-ahead equity risk premium, with the weights selected so that this single measure explains as much of the variability across models as possible (for the geeks: it is the first principal component). The value of 5.4 percent for December 2012 is about as high as it's ever been. The previous two peaks correspond to November 1974 and January 2009. Those were dicey times. By the end of 1974, we had just experienced the collapse of the Bretton Woods system and had a terrible case of stagflation. January 2009 is fresher in our memory. Following the collapse of Lehman Brothers and the upheaval in financial markets, the economy had just shed almost 600,000 jobs in one month and was in its deepest recession since the 1930s. It is difficult to argue that we're living in rosy times, but we are surely in better shape now than then.

Today's equity premium has reached a historic high

Percentage annualized




Sources: Authors' calculations; Barclays; Deutsche Bank; Duke/CFO Business Outlook survey; Federal Reserve Board; Federal Reserve Bank of New York; Goldman Sachs; J.P. Morgan; Nomura; the Center for Research in Security Prices; Federal Reserve Economic Data; Thomson Reuters; the websites of NYU's Aswath Damodaran; Dartmouth's Kenneth French, University of Lausanne's Amit Goyal, University of California at Berkeley's Martin Lettau, Yale's Robert Shiller.

The next chart shows a comparison between those two episodes and today. For 1974 and 2009, the green and red lines show that the equity risk premium was high at the one-month horizon, but was decreasing at longer and longer horizons. Market expectations were that at a four-year horizon the equity risk premium would return to its usual level (the black line displays the average levels over the last fifty years). In contrast, the blue line shows that the equity risk premium today is high irrespective of investment horizon.

The equity premium is elevated at all horizons

Percentage annualized




Sources: Authors' calculations; Barclays; Deutsche Bank; Duke/CFO Business Outlook survey; Federal Reserve Board; Federal Reserve Bank of New York; Goldman Sachs; J.P. Morgan; Nomura; the Center for Research in Security Prices; Federal Reserve Economic Data; Thomson Reuters; the websites of NYU's Aswath Damodaran; Dartmouth's Kenneth French, University of Lausanne's Amit Goyal, University of California at Berkeley's Martin Lettau, Yale's Robert Shiller.

Why is the equity premium so high right now? And why is it high at all horizons? There are two possible reasons: low discount rates (that is, low Treasury yields) and/or high current or future expected dividends. We can figure out which factor is more important by comparing the twenty-nine models with one another. This strategy works because some models emphasize changes in dividends, while others emphasize changes in risk-free rates. We find that the equity risk premium is high mainly due to exceptionally low Treasury yields at all foreseeable horizons. In contrast, the current level of dividends is roughly at its historical average and future dividends are expected to grow only modestly above average in the coming years.

In the next chart we show, in an admittedly crude way, the impact that low Treasury yields have on the equity risk premium. The blue and black lines reproduce the lines from the previous chart: the blue is today's equity risk premium at different horizons and the black is the average over the last fifty years. The new purple line is a counterfactual: it shows what the equity premium would be today if nominal Treasury yields were at their average historical levels instead of their current low levels. The figure makes clear that exceptionally low yields are more than enough to justify a risk premium that is highly elevated by historical standards.

The equity premium is high because Treasury yields are low

Percentage annualized




Sources: Authors' calculations; Barclays; Deutsche Bank; Duke/CFO Business Outlook survey; Federal Reserve Board; Federal Reserve Bank of New York; Goldman Sachs; J.P. Morgan; Nomura; the Center for Research in Security Prices; Federal Reserve Economic Data; Thomson Reuters; the websites of NYU's Aswath Damodaran; Dartmouth's Kenneth French, University of Lausanne's Amit Goyal, University of California at Berkeley's Martin Lettau, Yale's Robert Shiller.

But none of this analysis matters if excess returns are unpredictable because the equity risk premium is all about expected returns. So...are returns predictable? The jury is still out on this one, and the debate among academics and practitioners is alive and well. The simplest predictive method is to assume that future returns will be equal to the average of all past returns. It turns out that it is remarkably tricky to improve upon this simple method. However, with so many models at hand, we couldn't help but ask if any of them can, in fact, do better.

The table below gives the extra returns that investors could have earned by using the models instead of the historical mean to predict future returns. For investment horizons of one month, one year, and five years, we pick the best model in each of the four classes we consider together with the weighted average of all twenty-nine models. We compute these numbers by assuming that investors can allocate their wealth in stocks or bonds, and that they are not too risk-averse (for the geeks again, we solved a Merton portfolio problem in real time assuming that the coefficient of relative risk aversion is equal to one). The table shows positive extra returns for most of the models, especially at long horizons.

Model performance is varied, but better at longer horizons

Basis points per year



Sources: Authors' calculations; Barclays; Deutsche Bank; Duke/CFO Business Outlook survey; Federal Reserve Board; Federal Reserve Bank of New York; Goldman Sachs; J.P. Morgan; Nomura; the Center for Research in Security Prices; Federal Reserve Economic Data; Thomson Reuters; the websites of NYU's Aswath Damodaran; Dartmouth's Kenneth French, University of Lausanne's Amit Goyal, University of California at Berkeley's Martin Lettau, Yale's Robert Shiller.

Notes: We tested twenty-nine models in four classes (surveys, dividend-discount models, cross-sectional regressions, and time series regressions) over three investment horizons. In this chart, we plot the single best-performing model in each category. We also show how the optimal weighted average (the first principal component) of all models performs.

At face value, this result means that the models are actually helpful in forecasting returns. However, we should keep in mind some of the limitations of our analysis. First, we have not shown confidence intervals or error bars. In practice, those are quite large, so even if we could have earned extra returns by using the models, it may have been solely due to luck. Second, we have selected models that have performed well in the past, so there is some selection bias. And of course, past performance is no guarantee of future performance.

Disclaimer

The views expressed in this post are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.