

Forecasting the Equity Premium: Where We Stand Today

Matthew Spiegel

Yale School of Management

The Review of Financial Studies has among its missions the facilitation and promotion of a vigorous academic debate across unsettled questions in finance. This issue represents a cross section of views regarding one such debate: Can our empirical models accurately forecast the equity premium any better than the historical mean? Or, is the forecast our empirical models give us any more accurate than what we would get by simply using the historical mean?

The Goyal and Welch (2008) article that leads off this issue argues that the historical mean has done as well at forecasting the equity premium as any of the more complex empirical models that have been brought to bear on this issue. Given the numerous papers claiming otherwise this conclusion comes as something of a surprise. To some degree the discrepancy can be tracked down to how one handles the 1973–74 stock market collapse. Goyal and Welch show that the prescience attributed to the existing econometric models comes almost entirely from these two years of data. Since then the models have not performed better out-of-sample than the historical mean and did not do much better before then, either.

Goyal and Welch's (2008) working paper created enough of a controversy within the profession that papers by Campbell and Thompson (2008) and Cochrane (2008) soon followed with opposing views. Campbell and Thompson argue that the empirical models can yield useful out-of-sample forecasts if one restricts their parameters in economically justified ways. In contrast, Cochrane (2008) argues that the types of out-of-sample tests performed by Goyal and Welch are relatively weak, and that in-sample tests provide far greater power and can be convincing on their own. As the journal's executive editor, I do not wish to appear to take sides here (or, if you want to be less personal, you could say "The journal isn't taking sides; we leave it to our readers. . ."). Thus I am going to leave it to our readers to parse the various positions and determine for themselves which side has the better case.

The next two papers in this issue deal with the econometric issues involved with both determining if a particular set of results is meaningful and where the problems lie in attempting to forecast the equity risk premium. Boudoukh, Richardson, and Whitelaw (2008 [HB1]) start by noting that several existing studies claim that long-horizon return predictability is much better than

short-horizon predictability. However, they argue that this finding is largely an illusion and (or, to be more specific, “an illusion that arises”) arises simply from the mechanical growth in a model’s R^2 statistic with the estimated horizon length.

The special issue wraps up with an article by Lettau and Van Nieuwerburgh (2008). Their paper examines a regime-switching model in an attempt to see if explicitly allowing for structural changes in the economy can produce superior forecasts. While the answer to that particular question is “no,” the results do help to point to where the problems lie. They find that a regime-switching [HB1]model can determine if a structural break in the economy has happened shortly after it takes place. The problem is that the model apparently cannot determine what the new regime parameter values are quickly enough to produce useful real time forecasts (at least useful in the sense of providing trading information superior to what one would obtain by simply using the historical mean). This opens up an interesting question for future research: can academics produce an empirical model that allows for economic changes over time that is also capable of determining the “right” parameter values in time to help investors?

I hope you find the articles in this issue as interesting and thought provoking as I did. Here at *The Review of Financial Studies* we look forward to seeing where this debate leads us.

References

- Boudoukh, J., M. Richardson, and R. F. Whitelaw. The Myth of Long-Horizon Predictability. *The Review of Financial Studies* 21(4):1577–1605.
- 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* 21(4):1509–31.
- Cochrane, J. H. The Dog That Did Not Bark: A Defense of Return Predictability. *The Review of Financial Studies* 21(4):1533–75.
- Goyal, A., and I. Welch. A Comprehensive Look at the Empirical Performance of Equity Premium Prediction. *The Review of Financial Studies* 21(4):1455–1508.
- Lettau, M., and S. V. Nieuwerburgh. Reconciling the Return Predictability Evidence. *The Review of Financial Studies* 21(4):1607–52.

Copyright of Review of Financial Studies is the property of Society for Financial Studies and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.