Conditional Dynamics and the Multi-Horizon Risk-Return Trade-off
By Mikhail Chernov, Lars Lochstoer, and Stig Lundeby
Executive Summary

Many investors seek investment strategies that maximize long-horizon performance. A critical component of evaluating investment performance is assessment of risk exposure, which is measured using covariances. Because covariances are measured more precisely with higher-frequency data, both industry and academic practice is to risk-adjust returns at relatively short return horizons, typically monthly or quarterly. The classic example are the ubiquitous beta-pricing factor models and the associated return "alpha" (e.g., Gibbons, Ross, and Shanken, 1989, Fama and French, 1993).

However, short-run returns may not be good proxies for long-run returns, particularly if returns are not i.i.d., as pointed out in the early literature on tests of the (C)CAPM (e.g., Grossman, Melino, and Shiller, 1987, Longstaff, 1989). Related, while the existence of a monthly "alpha" means that one can improve the monthly Sharpe ratio by judiciously combining the asset at hand with the benchmark factors, monthly Sharpe ratio maximization is generally not the goal for long-horizon investors. In this paper, we ask two questions: How should one risk-adjust returns for long-horizon investors? And do our current state-of-the-art empirical risk models adequately account for longer-run risk-return trade-offs?

In our answer to first question, we use the stochastic discount factor to risk-adjust as this enables us to consider a given model's implications at different horizons. We show that risk-adjustment at longer horizons can be recast as a one-period risk-adjustment with managed portfolios (see Cochrane, 2005). In contrast to the literature, the multi-horizon return instrument is endogenous to the model at hand and involves lags of pricing errors at different horizons. The one-period moment conditions we derive are uncorrelated across time, which provides an efficient solution to an econometric issue that longer horizons means less independent data for a given sample.

To answer the second question, we ask if standard factor models, such as the Fama-French models, with and without a momentum factor, and more recent factor models of Daniel et al. (2018) and Stambaugh and Yuan (2017), properly account for risk at longer horizons. All of these models claim to do well at matching the cross-section of stock returns at the monthly horizon. Our tests are based on a simple idea. As an example, consider an investor who each month invests her wealth in the risk-free rate and the value factor (HML) with an investment horizon of, say, 24 months. If the Fama-French model is correct, the risk-adjusted returns to this strategy should be zero. However, we show that risk-adjusted returns appear large when aggregating this model to the 24-month horizon. In fact, we reject all the considered models when asking if they price such a simple trading strategy for each of the factors in each given model across multiple horizons.

The reason the models fail is that, while they price their own factors at the monthly horizon unconditionally, they do not price the factors conditionally. In particular, small persistent conditional pricing errors aggregate to large pricing errors over longer return horizons. We show explicitly how considering multi-horizon returns amount to a conditional test of the models.