The Information Cliff

A Brief Summary

Link to the full paper: https://yeli-macrofinance.com/infocliff.pdf

Understanding financial market participants' information about economic fundamentals is crucial: asset prices reflect expectations of future economic conditions, directly influencing resource allocation and policymaking. A large literature studies how information is processed, often emphasizing biases that distort expectation formation. In this paper, we begin by characterizing the supply of information and uncover an information cliff—a sharp decline in available information on economic fundamentals beyond the one-year horizon—and show that it sheds new light on central objects of canonical asset pricing frameworks—the dynamics of expected cash flows and expected returns.

The information cliff emerges from the horizon of firms' information supply to equity investors via corporate disclosure, particularly the quarterly guidance that firms release, and we find that more than 90% of guidance pertains to performance within the next year. The cliff is mirrored in the information content of analyst forecasts: forecasts of aggregate earnings growth for the upcoming year deliver an R^2 above 70% in predicting realized growth, but the R^2 drops significantly when forecasts for the second or third year are used to predict growth in the corresponding periods, and long-term growth forecasts are similarly weak. Moreover, since corporate announcements in the first month of a quarter are the most informative about aggregate economic conditions, the wedge between the predictive power of next-year forecasts and that of longer-horizon forecasts should be—and indeed is—most pronounced in those months. The alignment between the timing of information supply and forecast accuracy across horizons lends support to the one-year cliff.

To derive asset pricing implications, we extend the model of Lettau and Wachter (2007), allowing the dynamics of cash flows and risk prices to be driven by an arbitrary set of latent state variables and shocks, and we do not impose rational expectations. The first step is to identify the state-space dimensionality. We show that the logarithms of dividend strip prices scaled by realized dividends ("valuation ratios") are linear functions of the state variables, so the valuation ratios of strips across maturities span the state space. Computing the valuation ratios of S&P 500 dividend strips (Binsbergen, Brandt, and Koijen, 2012), we find that the first two principal components account for 96% of the total variance, and pairs of valuation ratios forecast returns and dividend growth as well as three or more do.

Therefore, a two-dimensional state space suffices, despite the many economic forces affecting the aggregate stock market. The two state variables can be rotated to represent the conditional expected return and the conditional expected dividend growth rate, each following an AR(1) process—a structure that prior studies have assumed but not tested.

The conditional expected dividend growth rate encapsulates the information that investors have about future cash flows. A non-zero autoregressive coefficient would mean that the current information set contains persistent signals, allowing participants to forecast growth two, three, and more years ahead. The information cliff therefore implies a zero autoregressive coefficient—a sharp restriction on the persistence of expected cash-flow growth, one of the central objects in the asset pricing literature. Using analyst forecasts to proxy for cash-flow expectations, we estimate this coefficient to be consistently around zero across specifications, and fitting a state-space model to dividend data with a latent expected growth rate yields the same conclusion.

We then connect the information cliff to the expected market return. In the model, if expected cash-flow growth has a zero autoregressive coefficient, the expected return is a univariate function of the slope of the valuation term structure—the wedge between the market price-dividend ratio (pd) and the valuation ratio of the one-year strip; otherwise, the sign of the slope's return forecasting error aligns with that of the autoregressive coefficient. The return predictive power of the slope is intuitive under the information cliff: because market participants have little information on cash flows beyond the next year, a steepening of the valuation term structure is not driven by improved long-term growth prospects but by a lower discount rate, which benefits the valuation of long-duration cash flows more; a steepening thus predicts lower market returns. Empirically, the slope delivers an in-sample R^2 of 25% and an out-of-sample R^2 of 15%, subsumes the predictive power of pd, and cannot be improved upon by including other predictors—in line with the expected return being a univariate function of the slope. A linear regression on the slope even outperforms machine learning algorithms that nonlinearly aggregate a large set of predictors (Kelly, Malamud, and Zhou, 2024). And in rolling windows, when the estimated autoregressive coefficient of expected cash-flow growth deviates from zero, its sign aligns with the slope's return forecasting error, exactly as the model predicts.

Beyond discount-rate variation, our results admit a mispricing interpretation that reorients market timing. Traditional timing bets against the overall valuation level, pd. Since market participants are well informed about the near term but face an information cliff beyond one year, mispricing is likely to reside in long-term cash flows, and the focus should be on the slope rather than the level: to time the market is to reduce exposure when the valuation term structure steepens and increase exposure when it flattens—betting against exuberance or pessimism about the long term. This strategy delivers a Sharpe ratio of 0.58, 55% higher than that of buy-and-hold.

References

Binsbergen, J. H. van, M. W. Brandt, and R. S. J. Koijen (2012). On the timing and pricing of dividends. The *American Economic Review* 102(4), 1596-1618.

Kelly, B., S. Malamud, and K. Zhou (2024). The virtue of complexity in return prediction. The *Journal of Finance* 79(1), 459-503.

Lettau, M. and J. A. Wachter (2007). Why is long-horizon equity less risky? A duration-based explanation of the value premium. The *Journal of Finance* 62(1), 55-92.