

Classes 5 & 6: The Cross-Section of Expected Returns

The materials to be covered in this class are among the most influential work in Finance since 1990s, with professor Eugene Fama as an influential leader. If you search online, there is a recorded talk by professor Fama on “A Brief History of the Efficient Market Hypothesis,” introduced by professor John Cochrane and took place at University of Chicago in 2008. Toward the end of the talk, he talks about his research philosophy.

What’s my research philosophy? Keep it simple. Keep it simple. Do things that people can understand. That’s easy for me because I am a simple-minded guy. If I can understand it, almost anybody can understand it. And my good friend David Booth told me twenty years ago that every business is 50% marketing. And academic business is not an exception. It is at least 50% marketing. The way you write has as big an impact on your impact in the world as a whole as what you actually do. So saying things clearly and simply is going to get you much further.

Around 2000, when I first became an assistant professor, this strand of literature was informally referred to as low-tech empirical work, while the more structural estimations in the vein of Hansen and Singleton (1982) were considered as high-tech, which I used to belong. In fact, my training was so narrow that I didn’t know much about the low-tech work. But once I understood what the cross-sectional approach is about, I was hooked. It was as if I was given a pair of new and powerful goggles that has color and is 3D, helping me see patterns in financial data that I didn’t know existed before. As empiricists, working with data is what we do all the time, and pure data work does not require a lot of techniques – anybody knows how to program can do it. And yet, keeping it simple is the hardest thing in academic research. It means that you can no longer hide behind fancy techniques or complicated models. For me, the high-tech person, the challenge was enormously appealing and I became a better empiricist because of that.

1 Fama and French (1992, 1993)

In the same talk mentioned above, professor Fama has this to say about Fama and French (1992):

Then we come to the time when Ken French comes here and he has the office next to me. And we start working together. And I figure, here is a guy who's got incredible energy and he can keep my reputation going for another twenty-five years. So I hooked up with him, incredibly smart guy, same work habits as I do, and we just work together very well. One of the big results was this paper that we wrote in 1992 called the cross-section of expected stock returns, which I never expected to get published. Because there was absolutely nothing new here. All we did was to take previously so-called anomalies and put them all together in one place. So what had happened? The CAPM came along in 1965 and people didn't start testing it until the early 1970s. There were papers by Fama and MacBeth (1973) and Black, Jensen, and Scholes (1973) and other people were testing the CAPM. The initial tests looked like they gave pretty good support to the model. Then in the early 1980s we started to come across things that are causing embarrassments for the model. But they came up one at a time. So people looked at them one at a time and said, one thing, that's not too bad. All we did was to put them all together and say, this thing just didn't work.

The Big Picture: I view Fama and French (1992) as a methodological paper, which provides guidance on how to extract information from a large cross-section of stocks with the objective to estimate and evaluate asset-pricing models. The recipe is amazingly simple: form portfolios by sorting stocks by characteristics, or pre-ranking beta's on the risk factors of interest. The sorting is dynamic: once a year for size, market beta, and book-to-market in Fama and French (1992). For more volatile variables, such as past returns in Jegadeesh and Titman (1993), the sorting is done once a month. Compared with using machine learning or other algorithms to look for patterns, the Fama-French approach is more appealing as the search is guided by economics and market intuitions.

If Fama and French (1992) is about sorting stocks into portfolios, then Fama and French (1993) is about using these portfolios to construct risk factors and test asset-pricing models. Relative to using individual stock returns directly, the firm-level idiosyncratic risks have been largely reduced in the portfolio approach. The presence of these portfolios also offers diagnostics on why the asset-pricing model fails. For example, focusing on the CAPM alphas of the size- and value-sorted portfolio, we learn that the CAPM beta fails to explain the large cross-sectional variation in expected returns of these portfolios, with value outperforming

growth and small outperforming big. For the investing community, such results have immediate implications, and for this reason, many of professor Fama's former students were early adopters of the quant approach, with some becoming very successful investors in the quant investing space. Today, Chicago's graduate school of business is named after David Booth, professor Fama's former research assistant and co-founder and executive chairman of Dimensional Fund Advisors.

Implications for the CAPM: As mentioned in professor Fama's talk, researchers started to test the CAPM in the early 1970s and the empirical evidence was supportive of the model. Post 1990s, with the publication of Fama and French (1992), there have been mounting empirical evidences against the CAPM. Focusing on the relation between expected return and beta, Fama and French (1992) show that the cross-sectional variation in expected stock returns is related to size not beta. Focusing on the CAPM alphas of the test portfolios, including the famous 25 Fama-French portfolios that are sorted by size and book-to-market, Fama and French (1993) reject the CAPM via the GRS test of Gibbons, Ross, and Shanken (1989). Interestingly, even the Fama-French three-factor model is rejected by the test.

It should be mentioned that, without the CAPM, none of the papers testing the CAPM would even be possible. Conceptually, with respect to the historical evolution of our collective understanding of financial risks, the CAPM plays a crucial role. It separates the systematic risk from the idiosyncratic risk, emphasizing the importance of this systematic risk in asset pricing. Empirically, the market portfolio remains a dominant presence. Pick any equity portfolios including the Fama-French portfolios, the market risk factor commands a decisive explanatory power, while the importance of the new risk factors such as HML and SMB is only secondary. From these perspectives, the CAPM has been an enormously successful model.

Post-Ranking Betas: Fama and French (1992) devote quite some space in discussing the pre-ranking and post-ranking betas. For any given risk factor, traded or non-traded, individual stock returns are regressed on the factor to obtain the pre-ranking betas. In Fama and French (1992), this regression is done once a year using a five-year rolling window of past monthly returns. Each year, stocks are sorted by their pre-ranking beta's into portfolios. Once the portfolios are formed, the pre-ranking betas are thrown out, and the post-ranking betas, estimated by regressing portfolio returns on the factor, are used for asset-pricing tests and estimation. In Fama and French (1992), the post-ranking betas have a very robust spread, indicating the persistence of CAPM beta at the individual stock level. For any asset-pricing tests using equity return data, this is the best scenario because the market portfolio is the most important risk factor in the equity market. For many other risk factors,

for example, the liquidity factor of Pastor and Stambaugh (2003) and the volatility factor of Ang, Hodrick, Xing, and Zhang (2006), constructing test portfolios with a wide enough spread in post-ranking betas turns out to be a big challenge.

Explaining SMB and HML: There are three strands of literature. Violations of the CAPM does not equate market inefficiency. For the rational camp, HML and SMB contain information above and beyond that in the market return for forecasting GDP growth. They could be driven by state variables that forecast time-varying investment opportunities or time-varying risk aversion. For example, the value premium might be driven by the premium for a systematic distress risk, and the size premium might be linked to a systematic illiquidity risk. The empirical evidences for these hypotheses, however, are mixed.

For the behavioral camp, the size and value anomalies could be associated with expectational errors made by investors. For example, in “Good News for Value Stocks: Further Evidence on Market Efficiency,” La Porta, Lakonishok, Shleifer and Vishny (1997) hypothesize that investors initially over-react to the negative information of value stocks by expecting them to continue slow growth. As they later adjust their expectational errors, the previously depressed prices of value stocks bounce back. Hence the superior returns associated with the value stocks. To test this story, they focus on the market’s reactions to earnings announcements of value and growth firms, and find that one-third of the HML returns is realized on the 3-day window around the earnings announcement day. Finally, the critics raise the issues of survival biases and data snooping.

2 Cross-Sectional Estimation of Factor Risk Premiums

The Fama-French approach of using sorted portfolios to test asset-pricing models can be extended more broadly to estimate the market price of any risk factors, traded or non-traded.

The Equity Risk Premium: The most straightforward estimation of the equity risk premium is via the difference between the expected return of the S&P 500 index ($\beta = 1$) and the riskfree rate ($\beta = 0$). Via the CAPM, the premium for the equity risk factor can also be estimated using the cross-section of expected stock returns. Fama and French (1992) provides guidance on how test portfolios can be constructed for this purpose. This estimation approach is driven mostly by intuition, and its econometric efficiency remains to be examined. Nevertheless, it has been widely adopted because of its simplicity. The failure of the CAPM indicates that the cross-section of stocks fail to price the market risk, although the

market portfolio as proxied by CRSP value-weighted index or the S&P 500 index demands a positive and statistically significant premium over the riskfree return.

Liquidity Risk Premium: Pastor and Stambaugh (2003) build a monthly time-series of aggregate liquidity measure from the stock-level liquidity measures, which capture the price impact of order flow for each stock. Unlike the risk factors in Fama and French (1993), this liquidity measure does not have a traded portfolio directly associated with it. The information contained in the cross-section of stocks becomes useful in gauge the asset-pricing impact of this liquidity measure. Instead of using the pre-ranking liquidity beta directly to sort stocks into test portfolios, Pastor and Stambaugh (2003) use predicted liquidity beta, which is a linear function of stock characteristics including firm-level historical liquidity beta, firm-level liquidity measure, trading volume, stock price, shares outstanding, past return and volatility. The post-ranking betas of the ten sorted portfolios have the desired spread, although the ranking is not monotonic and the spread is somewhat weak. For non-traded risk factors, this problem of having a robust spread in post ranking beta is quite pervasive. It is an indication that, unlike the market portfolio, such risk factors are not important in explaining the time-series variations of stock returns. Examining the difference in expected returns across the ten portfolios, they find that, from 1966 through 1999, the average return on stocks with high liquidity beta outperforms that for stocks with low liquidity beta by 7.5 percent per year. In other words, the aggregate liquidity measure is priced by the cross-section of stocks.

Volatility Risk: The first half of Ang, Hodrick, Xing, and Zhang (2006) focuses on the market price of volatility risk using the CBOE VIX as a proxy. They use pre-ranking VIX beta to sort stocks into five test portfolios and find that stocks with high VIX beta have low average returns. Given that worsening market condition is in general associated with sudden increases in VIX, this result makes perfect sense and this is also how the volatility risk is priced in the option market. The problem is that, the post-ranking betas of these test portfolios have a rather narrow spread. To show that the test portfolios exhibit high loadings with volatility risk over the same period used to compute the alphas, they use the test portfolios to construct a factor mimicking portfolio for the volatility risk and show that the post-ranking betas on this factor mimicking portfolio has a robust spread.

Funding Liquidity Risk: Hu, Pan and Wang (2013) propose a market-wide liquidity measure by exploiting the connection between the amount of arbitrage capital in the market and observed noise in U.S. Treasury bonds, using the intuition that shortage of arbitrage capital allows yields to deviate more freely from the curve and results in more noise in prices. We use the cross-section of hedge fund returns to estimate this liquidity risk premium. Sorting hedge

funds by their noise beta into 10 hedge-fund portfolios, we find that higher average returns for hedge funds with negative noise beta – those funds experiencing negative returns when the market-wide liquidity risk spikes up. This liquidity risk premium explains why some hedge funds can generate superior performance through exposures to a priced, market-wide liquidity risk factor. Interestingly, such highly exposed hedge funds are also found to have a higher exit rate in 2008 relative to the graveyard sample. Learning from the past papers, we pay special attention to the post-ranking noise beta to make sure that there is a direct link between the observed alpha and the liquidity risk exposure. We further use currency carry returns as an alternative set of test portfolios. we find that the asset currencies tend to experience negative returns when the liquidity risk spikes up, yielding a negative noise beta, while the funding currencies tend to provide positive returns when the liquidity risk spikes up, yielding a positive noise beta. Using the liquidity risk premium estimated from hedge fund returns to make risk adjustment, we find that the superior performance of the asset currencies is lower in magnitude and no longer statistically significant. In other words, high exposure to market-wide liquidity risk is a key driver for currency carry profits.

3 Cross-Sectional Patterns and Quant Investing

Quant Signals: Looking for patterns in the cross-sectional data has been an extremely active research area. It also has an immediate implication for the quant investing space. Popular quant signals include basic patterns such as size, value, and momentum. In addition, profitability, measured as the earnings-to-sales ratio, is also a useful quant signal. Recently, Fama and French propose a profitability signal that uses the operating profit (=revenues minus cost of goods sold, minus selling, general, and administrative expenses, minus interest expense) divided by book value of equity. They find that stocks with robust profitability overperforms stocks with weak profitability and create a factor called RMW (robust minus weak).

Accounting data plays an important role in signal creation. For example, Sloane (1996) shows that stock with low earnings quality (high accrual) underperform stocks with high earnings quality. The quant signal using analysts forecast revision also comes from the accounting literature. Stickel (1991) finds that analysts revision affect prices, but prices do not immediately assimilate the information. In fact, prices continue to drift in the direction of the revision for about six months after the revision. Another important pattern related to earnings news is reported by Bernard and Thomas (1989). This is the famous post earnings announcement drift: stocks with positive earnings surprises on their announcement day keep drift upward in their stock prices a few weeks (up to 60 days) after the announcement while

stocks with negative earnings surprises keep drifting downward.

Management impact has also been used as a quant signal. Loughran and Ritter (1994) reports long-term underperformance after IPO or SEO, and Ikenberry, Lakonishok, and Vermaelan (1995) reports long-term overperformance after announcements of share repurchases. In their recent paper, Fama and French introduces a signal that is similar in spirit. They use the firm's asset growth as a signal for firm investment, and find that stocks with low investment (low asset growth) outperform stocks with high investment. Calling firms with low investment conservative, and high investment aggressive, Fama and French introduce a new factor called CMA (conservative minus aggressive). Together with the market portfolio, SMB, HML, RMW (just mentioned), Fama and French build a new five-factor model.

Crowded Trades and Over-Used Signals: Popular quant signals are common knowledge, making quant investing an over-crowded space with over-used signals. The transparency of these trading strategies also makes the funds' portfolios and transactions easy to predict, inviting front runners. The 2007 quant meltdown is a result of this over-crowding, which is reminiscent the 1998 LTCM crisis, when the fixed income arbitrage space became over-crowded. In the case of LTCM, the actual trigger was Russia's default on its local currency debt, which LTCM did not have a lot of exposure to. Likewise, the initial trigger for the 2007 quant meltdown was disruptions in the sub-prime mortgage market, which most of the quant funds did not have any direct holdings.

What next? The search for new quant signals is still on. Given the massive amount of data mining in the past ten to twenty years, the number of interesting signals left for us to discover is shrinking, and this area is just not as exciting and creative as it was ten or twenty years ago. The next phase might be in the intersection of Finance and Technology, as investors are increasingly using alternative data and more sophisticated algorithms to look for tradeable patterns.