Summary of the Autumn Seminar

INQUIRE Europe

October 2-4, 2011

Le Royal Hotel, Luxembourg

Summarized by:
Maarten van Ravenswaaij
(AEGON Asset Management)

Table of Contents

Page number

Program of the Seminar

Monday, October 3rd, 2011

Key note address: Risk and returns after the crisis
Rene Stulz, Everett D. Reese Chair of Banking and Monetary Economics, The Ohio State University

The market premium for dynamic tail risk
Lorán Chollete, University of Stavanger

Extreme dependence structures and the cross section of expected returns
Stefan Ruenzi University of Mannheim

An asset pricing approach to liquidity effects in corporate bond markets
Dion Bongaerts, Erasmus University

Private equity performance and liquidty risk
Ludovic Phalippou, Oxford University

Liquidity risk and mutual fund returns
Ronnie Sadka, Boston College

Tuesday, October 4th, 2011

Should investors include commodities in their portfolios after all? New evidence
George Skiadopoulos, University of Piraeus

Tutorial on developments in behavioural finance
Kent Daniel, Columbia University

Stock return predictability and variance risk premia: statistical inference and international evidence
Tim Bollerslev, Duke University

Key note address: Risk and returns after the crisis
Rene Stulz, The Ohio State University

The last decades featured the financial crisis periods of 1987, 1998 and 2001-2002. If lessons from these periods were learned would the current crisis (2007-?) have occurred? It seems people think this crisis is different from the previous ones ("this time is different") but is it? Regression results from a paper by the presenter and others indicate that bank stock returns from the 1998 crisis explain bank stock returns during the recent crisis. So perhaps the banks were and are still exposed to the same systemic risk factor. After the 1998 crisis VaR was heavily criticized as a risk measure and proposed improvements such as extreme value theory and liquidity adjusted risk measures gained popularity. Yet, financial institutions that did poorly in 1998 did poorly in the recent crisis. Did we fail to learn, did we forget or was it not worth it to change?

In crisis times diversification breaks down as correlations go up. Even hedge fund returns are correlated in crisis times across hedge fund styles. High volatility is a possible partial explanation of correlation breakdown and fat tail return distributions. An additional explanation is the disappearance of liquidity as investors run to safe assets. Before the crisis a AAA rated structured finance security was viewed as riskless but when it became clear they were not, a reassessment of the risk took place and uncertainty caused liquidity to dry up for a while. Some markets stopped working and obvious arbitrage opportunities (interest rate parity, CDS-corporate bond spread basis) could not be exploited due to lack of liquidity to trade and due to lack of liquidity to fund the trade. Even worse, shocks to asset liquidity cause shocks to funding liquidity and vice versa. The interactions between positions, transactions, prices, leverage and margins typically cause a negative spiral and deteriorate a crisis situation. Especially banks which were very sensitive to funding liquidity shocks lost a lot during the crisis. Banks with more equity did better.

Should investors be compensated for liquidity risk? One has to distinguish between illiquidity and liquidity risk. This has implications for performance attribution. Is there an outperformance or is it just a compensation for illiquidity or liquidity risk? An example is LTCM which had a lot of exposure to liquidity risk. In this seminar the paper by Bongaerts, De Jong and Driessen reports a compensation for illiquidity but not for the liquidity risk of bonds. The paper by Franzoni, Nowak and Phalippou finds that private equity has no alpha when liquidity risk is taken into account. For mutual funds, Dong, Feng and Sadka find that mutual funds loading on liquidity risk perform better, but the better performance is only partially explained by the liquidity risk premium.

Two papers in the conference measure the crash exposure using approaches inspired by extreme value theory (EVT). Both papers identify a premium for exposure to their measure of tail risk. Interesting questions regarding this premium are: does that premium proxy time varying volatility and is there a non-linearity or systemic risk that beta does not capture?

Two practical lessons of the crises are: the importance of running liquidity dry-up scenarios and measuring exposure to tail risk. An investor such as a pension fund with low funding liquidity risk can improve performance by taking on liquidity risk. For a hedge fund with monhtly redemption taking on liquidity risk can have very adverse consequences.

The market premium for dynamic tail risk

Lorán Chollete, University of Stavanger

The paper presented by Lorán Chollete is joint work with Ching-Chih Lu. In this paper a measure of systemic risk is presented and estimated. The testing framework uses a stochastic discount factor model to price risky asset payoffs. The stochastic discount factor relates to the marginal utility of having a dollar today instead of later. The authors expect that in times of crisis investors require stocks that are exposed to tail risk to earn higher returns (the tail risk premium is positive).

According to extreme value theory the so called tail index is a useful measure of the heaviness of asset return tails. One shape parameter governs the tail behavior of the distribution and this parameter can easily be estimated using daily returns. For each day in the 1964-2009 period a single aggregate market tail index is constructed using returns of more than 2000 U.S. stocks. The correlation between the tail index and the market implied volatility index is about -0.2.

The exposure of annual stock returns to the tail index risk factor is estimated in a regression framework with the standard Fama-French risk factors and the liquidity measure of Amihud (2002). The correlation between the tail risk factor and the other variables is below 0.12 which suggests this factor captures additional information beyond the existing factors.

Subsequently stocks are sorted into quintiles according to tail risk exposure. A portfolio that goes long in stocks in the most exposed tail risk quintile and shorts stocks in the least exposed tail risk quintile earns 5.7% on an annual basis which indicates exposure to tail risk is rewarded with higher returns. The estimated risk premium of exposure to tail risk is between 1% and 3%. Results from standard asset pricing tests (the J-test, the Hansen-Jagannathan test and the delta-J test) confirm that the best model requires the tail risk factor to complement the existing risk factors. In the cross section, financial firms and firms with high leverage do not seem to be more exposed to the tail risk factor than other firms.

Extreme dependence structures and the cross section of expected returns
Stefan Ruenzi, University of Mannheim

The presentation on the crash sensitivity of stocks by Stefan Ruenzi is joint work with Florian Weigert. Evidence of crash aversion from the option pricing literature is the observed stylized fact that deep out of the money put options tend to be very expensive. Whether this effect is present in the cross-section of stock returns is examined in this paper. If so, the linear correlation in traditional asset pricing models cannot account for these dependencies.

The crash sensitivity of individual stocks is captured by the lower tail dependence (LTD) which measures the probability that a stock has a very negative return given that the market return is also very negative. Statistically it is possible for two stocks to have the same linear correlation with the market return but have a different degree of tail dependence. The main hypothesis to be tested is whether crash averse investors are compensated for holding stocks that exhibit a higher degree of lower tail dependence.

The dataset consists of daily returns of all U.S. common shares trading on the NYSE or AMEX in the period of 1963-2009. For each stock in each year the LTD is measured. As the yearly intervals used are non-overlapping, the contemporaneous relation between realized average returns and realized LTD's can be examined. Copula functions are used to capture the dependence structure of returns. For each stock and the market various parametric copulas are estimated and the copula that fits best is used to calculate the value of the LTD for a particular stock in each year.

The aggregate (value weighted) LTD fluctuates over time and the extreme values coincide with market crashes such as Black Monday (1987) and the Lehman crisis (2008). In order to examine the cross-sectional sample characteristics the stocks are assigned to quintiles according to their realized LTD value. The highest LTD quintile outperforms the lowest LTD quintile by 17%. These crash sensitive stocks also typically have higher betas, higher downside betas, tend to be stocks from larger companies and typically are growth stocks. After controlling for beta, the excess return of crash sensitive stocks remains about 12%. Similar excess returns are found when controlling for downside beta and co-skewness.

Apart from the bivariate results the impact of LTD is tested in a multivariate model. The LTD factor turns out to be a very significant determinant of returns, ranking third after the effects of beta and size. An increase of one standard deviation in LTD is associated with an average return premium of about 5% per year. The impact of downside beta vanishes once LTD is included in the regression model. The effect of LTD is very stable over time and during financial crises, high LTD stocks indeed perform significantly worse than low LTD stocks. Of all days in the sample where the market return performed worse than -5% the high LTD stocks on average had a return that was 5% worse than the low LTD stocks.

Finally, a trading strategy is evaluated that uses past (observable) LTD realizations. High past LTD stocks are bought and low past LTD stocks are sold. Gross of transaction costs this strategy earns about 0.3% excess return on a monthly basis after correcting for the usual risk factors. If the LTD risk factor would be added as a control variable that alpha may disappear. Overall there is empirical support that investors require a compensation for holding stocks with a strong sensitivity to extreme market downturns.

An asset pricing approach to liquidity effects in corporate bond markets
Dion Bongaerts, Erasmus University

The presentation by Dion Bongaerts is joint work with Frank de Jong and Joost Driessen. In the early years of this millennium the credit spread puzzle was an observed phenomenon: observed credit spreads were found to be much higher than justified by historical losses and reasonable risk aversion assumptions. In a sample the authors find a 1.2% spread for a typical AA bond while the expected default loss generates a credit spread of only 0.06%. In the literature additional risk factors were tested but the puzzle remained. The contribution of this paper is to examine whether differences in liquidity levels (expected liquidity) or liquidity risks can explain this puzzle using a formal asset pricing model.

Expected liquidity is an expected transaction cost that depends on the turnover of the portfolio. Liquidity risk is a risk factor, influenced by risk aversion, and depends on the timing and variation of liquidity costs. In the pricing model, the corporate bond returns are explained by the following systematic risk factors: equity market return, changes in corporate bond market liquidity, changes in equity market liquidity and changes in equity volatility. The expected corporate bond return is explained by the exposure to the risk factors multiplied by the risk premia of the factors plus an additional pricing effect of the expected liquidity level. The latter is measured by average transaction costs over the sample. In the estimation procedure the risk factor exposures are estimated first from the realized returns. In the second stage the expected returns are regressed on the factor exposures and the expected liquidity costs.

All trades in U.S. corporate bonds from October 2004 to December 2008 were used from the TRACE database. Bond portfolios are constructed based on credit risk ratings and liquidity proxies. For each portfolio the returns, the expected return and transaction costs are estimated. The expected return is estimated using yields and expected default losses from Moody's KMV. As bonds trade on an irregular basis the transaction costs are difficult to estimate. The implied bid-ask spreads and implied realized returns are estimated using Bayesian Gibbs sampling. This exploits the Roll (1984) estimator and assumes all bonds in the same credit risk / liquidity category have the same transaction cost.

The signs of the risk factors from the first stage of the estimation procedure as expected: corporate bond returns have positive exposure to stock market returns and negative exposures to liquidity and volatility shocks. Results from the second stage multivariate regression indicate that expected liquidity is a very significant determinant of expected returns but liquidity risk does not seem to matter. In fact, except for liquidity risk all other risk factors including equity market liquidity risk are statistically significant. The authors explain the irrelevance of liquidity risk by stating that liquidity risk only matters if investors would trade these assets when costs are high but in practice investors avoid trading these expensive assets and trade liquid assets instead. Several robustness checks were performed but these do not change the main results.

Private equity performance and liquidty risk
Ludovic Phalippou, Oxford University

The presentation by Ludovic Phalippou on private equity performance is joint work with Francesco Franzoni and Eric Nowak. Liquidity risk arises when investors have to pay higher transaction costs in bad market times compared to normal market circumstances. In this paper the diversification benefits of investing in private equity are examined. The exposure to risk factors are assessed, in particular the exposure to liquidity risk. Liquidity risk is not to be confused with the (generally smaller) liquidity level of this asset class.

A problem with private equity data is the absence of regular return observations. Cash flows occur infrequently. Several technical approaches have been put forward in the literature to deal with this issue. The innovative approach used in this paper assumes there is a single beta for all investments within the private equity asset class. This is sufficient to allow for the return model to be estimated. Cash flow data on more than 4400 liquidated private equity investments is used, covering a period of 20 years. This database is not publicly available. Roughly 50% of the firms is based in the U.S. and the other half in Europe.

Over the 1975-2007 sample period, the estimated gross rate of return is about 19%, net of fees it is about 12%. When moving from a CAPM to the 4 factor Pastor-Stambaugh return model, the excess return of private equity decreases to about 0.2% as private equity returns appear to be very correlated to the value premium (the Fama-French HML factor). The market beta of private equity is about 1.3, which is roughly in line with the cross sectional average leverage of the private equity firms. The exposure to the (public equity market) liquidity risk factor is positive and statistically significant as well. Within sample the 24% the cost of capital of the asset class is decomposed into 6% for the risk-free rate and 18% for the total risk premium. The latter is further decomposed into 10% for the market risk premium, 5% for the value premium and 3% for the liquidity risk premium and an excess return which is close to zero.

What is the economic explanation of these results? Further inspection at the investment level yields the observation that investments held during the best liquidity periods had returns above 20% whereas investments held during the worst liquidity liquidity periods had negative returns. A possible link between private equity returns and market liquidity occurs via the funding liquidity channel: times of low market liquidity coincide with times when private equity managers find it difficult to refinance their investments. In these periods they may be forced to liquidate the investments or to accept higher borrowing costs, which hurts returns. Data on changes in credit conditions (loosening/tightening) by U.S. banks from the Senior Loan Officer Surveys indicatively support this hypothesis. Consistent with this evidence is the finding that more dividends are paid when there are positive liquidity shocks and improvements in credit standards.

Liquidity risk and mutual fund returns
Ronnie Sadka, Boston College

The presentation on liquidity risk by Ronnie Sadka is joint work with Xi Dong and Shu Feng. The systematic liquidity risk, or liquidity beta, of several asset classes has been studied in the literature. This paper studies the liquidity risk of mutual funds. The liquidity factor of Sadka (2006,2010) is computed based on intraday price impacts. For each stock one observation of the liquidity factor is constructed for each month. Changes in this liquidity factor measure the liquidity risk.

The CRSP dataset contains almost 14000 mutual funds in the 1984-2009 period. A fund's liquidity risk is estimated in a regression of returns on the market return and the liquidity risk factor. Subsequently, portfolios of mutual funds are constructed using the returns of the past 12 months and the exposures to the Fama-French risk factors, momentum and the liquidity risk factor. The performance of these portfolios is measured in the next month. This exercise is repeated for each month. The excess return of going long in the most exposed decile of funds and short in the least exposed decile of funds is calculated for each of these risk factors and then corrected for exposure to the risk factors. Of all the 5 factors, only the excess return of the liquidity beta sorted portfolio is significant and it amounts to an annual performance of 6%. In other words, high liquidity beta mutual funds outperform low liquidity beta funds in the next period.

Is the observed return spread a reflection of a risk premium or skill? In order to investigate this the liquidity risk premium factor of stocks is constructed in the same way the liquidity risk premium factor of mutual funds was constructed and is used as an additional explanatory variable in a regression of liquidity beta return spreads on risk factors. It turns out the liquidity risk of stocks is a significant determinant of the liquidity beta return spread, but 80% of the return spread is still left unexplained. So 20% of the liquidity beta return spread is attributed to exposure to the liquidity risk premium and the rest to managerial skill. This result is similar across mutual fund investment styles. Evidence from holdings data supports these findings as the difference in holdings in high liquidity stocks by high liquidity beta mutual funds compared to low liquidity beta funds does not justify the excess return difference between these funds.

The causality relations between the liquidity return spread effect and other mutual fund stylized facts are examined. These stylized facts are the performance persistence effect, the fund size effect and the smart money effect. The liquidity risk effect is not explained by the other effects and the other effects are mainly present in the high liquidity beta quintiles.

In a previous paper a similar liquidity beta spread for hedge funds was found. Lead-lag analysis using a Vector Autoregression shows that the mutual fund liquidity spread leads that of hedge funds.

Should investors include commodities in their portfolios after all? New evidence
George Skiadopoulos, University of Piraeus

The presentation on commodities is joint work with Charoula Daskalaki. The paper studies whether investors are better off by adding commodities to a portfolio that is invested in cash, stocks and bonds. It addresses some of the shortcomings encountered in the large literature on the diversification benefits of commodities. For example some of the existing work uses a mean-variance setting but as higher moments may matter, this can be improved upon. Furthermore, some of the earlier work uses in sample analysis and does not consider the use of enhanced commodity indices.

The dataset consists of monthly closing prices from Bloomberg over the 1989 to 2009 period. The S&P500 total return index is used for the equities, the Barclays U.S. Aggregate Bond index for bonds and the LIBOR one-month rate for the risk-free rate. For commodities two indices are used (S&P GSCI and DJ-UBS CI) as well as individual futures contracts on cotton, copper, crude oil, gold and live cattle.

If the returns on commodities can be replicated using the other available asset classes, then this asset is spanned by these benchmark assets and it is not interesting for an investor to add commodities to the opportunity set. This hypothesis can be tested using regression analysis for both the mean-variance case and the non mean-variance case (which takes higher moments into account). For the mean-variance case there is spanning i.e. the introduction of commodities does not make the investor better off. In the non mean-variance case, there is no spanning i.e. commodities should be included. These results are in sample.

In the out of sample exercise the expected utility of an investor is maximized using historical samples. The optimal portfolio is held for one month and then optimized again. The sample sizes range from 36 months to 72 months. The performance of the optimal portfolios are evaluated using the Sharpe ratio, the opportunity cost of not investing in commodities, the portfolio turnover and a risk-adjusted measure net of transaction costs.

The out of sample results indicate that all performance measures deteriorate once commodities are included, with the exception of gold and during the commodity boom period of 2005 until June 2008. So in contrast to findings from earlier papers, according to these results commodities should not be included in investor’s portfolios. The same results are found for alternative specifications of the utility function and for second generation commodity indices that do not roll over to the shortest contract available but to the next contract to capture a positive roll return.

Tutorial on developments in behavioural finance

Kent Daniel, Columbia University

Behavioral finance attempts to explain financial market behavior such as asset pricing anomalies and corporate financing decisions with results from experimental psychology on e.g. information processing biases and irrational behavior. The latter means people use information imperfectly. Experimental psychology documents these regular “mistakes” or biases. The work of Kahneman and Tversky is especially well known in this field.

The efficient market hypothesis (EMH) as laid out by Fama in 1970 serves as an important baseline in understanding the biases observed in behavioral finance. Loosely it is based on the idea that investors see all available information and process it perfectly. As a result information and prices are directly linked and investor behavior can be ignored. The EMH implies a security’s price is always equal to the discounted value of expected future cash flows and changes only and immediately in response to new information about cash flows or discount rates.

Initial empirical tests provided strong support for the EMH. However in the last decades new evidence has revealed predictability seemingly inconsistent with the EMH such as the size effect, the book to market effect, the (price and earnings) momentum effect, the announcement effect, the accrual effect, the issuance effect and the idiosyncratic volatility effect. The existance of these anomalies received criticism in the literature such as being caused by data mining, observed without consistent patterns, or the result of improperly measured risk. In turn these criticisms received responses that argued for the presence of the anomalies as the effects were observed out of sample, repeatedly, across regions and asset classes and trading strategies based on the anomalies improve the risk corrected return (such as the Sharpe ratio).

From an academic perspective the key issue is understanding the source of the return differentials just described. From a portfolio management perspective a better understanding of investor “errors” should allow to avoid such errors and to develop trading strategies to profit from the errors of others. Evidence from experimental psychology suggests that the main errors, or biases that are observed are overconfidence (overestimating personal ability), anchoring (using seemingly irrelevant information for decision making), loss aversion and framing (people respond differently to losses and gains). A cognitive reflection task administered to diverse audiences including professional fund managers and students from MIT, shows that most people are prone to use their intuition when making decisions and behave in line with these biases.

A number of economists have criticized the the field of experimental psychology stating that it is just a bunch of results that lacks an underlying theory that drives the findings. A new field called neurofinance aims to close this gap by studying brain activity when individuals make choices or answer questions. One of the important techniques for doing this is functional magnetic resonance imaging which allows researchers to see which parts of the brain “light up” in the subjects. One study in particular investigates distinct areas of the brain associated with gains and losses.

Behavioral finance theories use the judgemental biases to explain the observed financial market anomalies. In a particular study by the presenter and co-authors they argue overconfidence is consistent with documented return patterns, including the value and momentum effects. For example, investors form estimates of future asset values (views) and overconfidence in these views causes them to slowly update to new information which generates a momentum effect. The implications of their model are also consistent with the size effect, the value effect, long horizon reversal and underreaction. Based on these arguments, Jiang, Lee and Zhang (2005) find empirical support for this model in a study of the momentum effect across a sample of firms with different levels of information uncertainty. For firms with larger information uncertainty, the larger momentum effect is consistent with the theory of slower updating of investor views.

Some challenges for the field of behavioral finance remain. One issue is assessing the magnitudes of the effects as currently most studies only address the direction of the effects. Another challenge is studying why and how arbitrageurs have been able to cope with some anomalies which seem to have disappeared or diminished and why other anomalies have persisted.

Stock return predictability and variance risk premia: statistical inference and international evidence
Tim Bollerslev, Duke University

The presentation by Tim Bollerslev is joint work with James Marrone, Lai Xu and Hao Zhou. It is a well known stylized fact that volatility clusters in time. The literature devoted to measuring, modeling, forecasting and pricing volatility can be divided in two streams. The first is concerned with direct volatility estimates and includes model-based estimation (ARCH, GARCH, stochastic volatility) and model-free (realized volatilities) estimation approaches. The second stream is concerned with implied volatility estimates and also includes model-based estimation (Black-Scholes, CIR, Heston) and model-free estimation approaches.

Realized volatility measures use the actual observed price data. Regardless of the model assumptions it can be show that the true (integrated) volatility can be approximated by taking the limit of the realized volatility while increasing the observation frequency. So the use of high-frequency data is crucial.

Implied volatility measures use derivative price data. Black-Scholes-Merton type models assume a constant volatility which is not consistent with e.g. the observed volatility smile in option markets. Later in time model-free implied volatility formulas were derived which use a weighting of available option prices. The current VIX and many other volatility indices are based on this idea. Empirically the option implied VIX typically exceeds the realized volatility, in other words, volatility risk is priced. This makes it profitable to sell volatility most of the time. This variance risk premium (VRP) can be seen as a proxy for aggregate risk aversion or as a proxy aggregate economic uncertainty.

Regression results show that VRP helps to predict monthly U.S. equity returns especially at the 3 month horizon. Other traditional predictor variables as well as realized volatility and implied volatility in isolation have no predictive power. Simulation results support the reliability of the findings. The study was reproduced for Germany, France, Japan, Switzerland and the U.K. for the more recent period of 2000-2010. As expected the VRP is positive for each country. Monthly returns and VRP's are highly correlated across countries and display a large negative spike at the start of the 2008 crisis. The return predictability pattern of the VRP of the countries shows similarities in shape but is statistically not as significant as the U.S. pattern.

The similarity in the pattern raises the question whether a global effect is present. Results from country specific regressions show the shape similarity of the return predictability patterns across countries. The return predictability of VRP is short lived and strongest at the 3 to 4 month horizon for all countries. Restricting the coefficients in a panel regression allows for even sharper statistical inference. The larger t-statistics of the panel regression coefficients support the idea of a global VRP rather than a country specific VRP. Temporal variation in aggregate risk aversion, aggregate uncertainty and the fear of tail events are important in order to explain the global VRP, but more research in this area is needed.