Journal Of Investment Management, Vol. 6, No. 2, (2008), pp. 1–7

www.joim.com

JOIM CONFERENCE SERIES-FALL 2007

1 Paul Samuelson

In his keynote address Paul Samuelson discusses the double-edge sword that is financial engineering. Specifically, he asks the pertinent question of whether the Surgeon General should declare that financial engineering is hazardous to society's health. Advances in financial engineering made by such great names as Louis Bachelier, Holbrook Working, Alfred Cowles III, Fischer Black, Myron Scholes, Robert Merton, and others have allowed risk-bearing to be more efficiently allocated to people best able to handle it and have created mechanisms to insure against risks. However, in the process many exotic financial instruments have been created in their wake, which are used for speculative purposes as opposed to reducing risk. For instance, financial engineering developed, with John Bogel, index funds, which are no-load, diversified, and tax-managed, for regular investors to take on a passive investment strategy. On the flip side, however, financial engineering also developed Exchange Traded Funds, in which inadequately informed investors have the temptation to actively speculate on the indices. Paul Samuelson draws a parallel with general scientific progress. The paradox of scientific progress is that on one hand it has allowed for improvements in quality of life and human longevity, but on the other hand it has amplified the harm that one or a few people can do. Take for example the case of a demented farm boy. In the old days, the worse he could do was torture his livestock. Today, he can go to a library and get the knowledge to build a nitrate bomb from his fertilizers to bring down a twenty-story building. Thus, financial engineering like scientific progress in general can have both good and bad consequences. The current worry is that a combination of high financial leverage, poor transparency, and cockeyed regulation is creating systemic risks in the financial markets. The answer to the problem is better regulation, because the cure for bad regulation is not zero regulation, as Milton Friedman suggested, but one whose good effects outweigh its negative consequences.

2 Robert Engle

Robert Engle in his keynote address examines the issue of whether we can anticipate future correlations in the equity markets, how and why correlations change over time and what are the models that can get to the best estimates of correlations for financial decision making. Correlations are important for several reasons. First, they are important in calculating portfolio risk. Second, they are necessary to form optimal portfolios. And third, they are a key input in pricing, hedging, and trading derivatives.

Correlations are time-varying as evidenced by the derivative prices of correlation sensitive products and today there are traded derivative products based on correlations. Correlations between equities change for several reasons. One reason is that companies change their lines of business becoming more or less like others. For instance, if one were to look at the correlation between American Express and General Electric, they would observe a growing correlation among the two. This is likely due to the fact that over the years the commercial and consumer finance divisions of General Electric have been accounting for a growing share of the company's total revenue. Thus, General Electric business has become more similar to American Express' business. Another reason for changing correlations is for the increased volatility in a common factor. For instance, rising energy prices and increased volatility in the energy markets have increased the correlation between Boeing and General Motors.

There are several models that can be used to estimate correlations for financial decision-making. There is the Dynamic Conditional Correlation (DCC) model, a one or multiple factor ARCH model, a Multivariate GARCH model, a factor DCC model, and a Dynamic Equicorrelation model. The DCC model is a new type of multivariate GARCH model that is particularly convenient for big systems. Essentially, to implement the DCC method, one needs to estimate volatilities for each asset and compute the standardized residuals or volatility-adjusted returns. The next step is to estimate the time-varying covariance between these assets using a maximum likelihood criterion and one of several models for the correlations. The third step is to form the correlation matrix which is guaranteed to be positive definite. Robert Engle shows how the DCC model can be applied to examine the correlation between American Express and General Electric. The DCC correlation estimates are very close to the 100-day historical correlation. Furthermore, he shows how DCC can be used to look at the correlation between the Shanghai Exchange A-Shares Index and the MSCI China Index. The MSCI China Index is available to global investors whereas A-Shares are available only to Chinese investors. The findings are that the volatility of the A-Shares is greater than the internationally traded China Shares and the correlation between these is small but is increasing as the Chinese government is opening up its local markets.

Factor models are also important to anticipating correlations because they allow factors to appear and disappear in the correlation model. There are several types of factor models. The One Factor ARCH model of Engle, Ng, and Rothschild, the One Factor GARCH model, and finally, the One Factor GARCH with DCC estimated correlations between all residuals called a Factor DCC model. In general, these factor models deliver superior performance.

3 Marti Subrahmanyam

In his presentation, Marti Subrahmanyam investigates whether a liquidity effect can explain the total yield spread on risky corporate bonds relative to their risk-less benchmarks after default risk has been accounted for. In order to gage the inherent default risk of the bond the study uses the spreads on Credit Default Swaps (CDS), which is then used to calculate the CDS-bond basis. The CDS-bond basis is the difference between the CDS spreads of the issuer and the par-equivalent corporate bond yield spread. The basis in this context can be interpreted as the non-default component of the yield spread and the study investigates whether bond liquidity can explain it. To measure bond liquidity the study uses a recent metric in the literature called latent liquidity. In simple terms, latent liquidity is a measure of how frequently the holders of a given bond turnover their portfolio. The argument is that a bond is likely to be more accessible if it is mostly held by investors who trade often. Information used in the computation of the latent liquidity was obtained from the State Street Corporation one of the largest custodians in the global financial markets. Information used in the computation of the CDS-bond basis was obtained from the Trade Reporting and Compliance Engine (TRACE) for bonds and the GFI Group Inc. and CMA DataVision for CDS. The final sample consists of a total of 33,000 bond-quarters from July 2002 to June 2006 on 4,972 bonds from 1,167 unique firms. The results of the study are, first, that liquidity explains a statistically and economically significant part of the CDS-bond basis and the bond yield. More liquid bonds have a lower CDS-bond basis and command higher prices. Furthermore, latent liquidity at the beginning of the quarter can be used to predict the basis during the quarter. The second finding is that the liquidity present in the CDS market itself affects the CDS-bond basis. The CDS-bond basis reflects a premium for the relative illiquidity of the corporate bond compared with that of the CDS. The less liquid the bond is relative to the CDS the more negative is the basis. The third main finding is that firm level variables that account for credit risk such as leverage and the proportion of tangible assets are significant in explaining the basis suggesting that the CDS price does not fully control for the credit risk of the bond. Finally, bond-level characteristics such as tax status and the presence or absence of covenants are important determinants of the CDS-bond basis. These findings suggest the presence of frictions in the arbitrage mechanism between the CDS and bond markets.

4 John Campbell

In his presentation John Campbell specifies and estimates a model of the nominal term structure of interest rates, which accounts for the time-varying correlation between inflation and real shocks. This model, which is outside of the affine-class of term structure model is driven by four state variables: the real interest rate, risk aversion, expected inflation and the covariance between nominal variables and the real economy. The model which is linear-quadratic is solved using a general result on the expected value of the exponential of a non-central chi-squared distribution and is estimated using a Kalman filter approach. Explaining the changing risks of nominal bonds is important for answering the question of whether nominal bonds are risky investments or hedging instruments. This question is complicated by the fact that the answer to this question is sensitive on the time period studied. For example, the covariance of nominal bond returns with stock returns in the 1970's and early 1980's was positive implying that bonds had a high Beta and were thus risky. However, in the late 1990's and early 2000's bond and stock returns were negatively related, implying a negative Beta on bonds and the possibility of using them to hedge shocks to aggregate wealth. Thus understanding the term structure of nominal interest rates requires an understanding of the changing covariance between nominal and real variables. The term structure model developed in this study implies that the risk premia of nominal bonds have changed over the years with movements in risk aversion, proxied by the equity dividend yield, and changes in the covariance between inflation and the real economy. Nominal bond risk premia were high in the early 1980's when bonds covaried with stocks and risk aversion was high. In 2000's when bonds and stocks moved in opposite directions the risk aversion was relatively low so negative bond risk premia were modest. Furthermore, the model can explain the finding by Cochrane and Piazzesi (2005) that a tent-shaped linear combination of forward rates predicts excess bond returns at all maturities better than maturity-specific yield spreads. In this model, the covariance between inflation and the real economy has opposing effects on longer-term bond yields. On one hand it raises yields by increasing the risk premium but on the other hand it lowers them through increased volatility. At the long end of the yield curve these two effects cancel for high levels of the nominal-real covariance, whereas at the intermediate portion of the curve the risk premium effect dominates. Therefore, the level of intermediate yields relative to short- and long-term yields is a good proxy for the nominal-real covariance and explains the risk premium on nominal bonds.

5 Andre F. Perold

Andre Perold in his presentation discussed risk stabilization and asset allocation. The most common asset allocation among pension funds, endowments, and wealthy individuals is a static policy allocation where the weights of different asset classes are kept constant until there is a policy change. The problem with this allocation strategy is that it assumes that the assets composing the portfolio have constant risk and constant covariance with each other. For example, a simple look at the time-series of the VIX index, which measures the implied volatility of the S&P 500, shows that equities exhibit time-varying risks. Furthermore, the correlation between equities and bonds is also time-varying, with negative correlation in the 1950's and 1960's, positive correlation from the 1970's to the 1980's, and then again a negative correlation in the late 1990's and 2000's. Thus, portfolios with static policy asset allocations are themselves going to have changing risk characteristics. A better asset allocation strategy might simply be to keep the risk of the portfolio constant as opposed to keeping asset weights constant. Andre Perold develops a simple model that can be used to evaluate both the static policy allocation and the constant risk allocation with respect to an optimal allocation strategy. The model assumes that allocation needs to take place between a risky asset and risk-less asset. The risk premium on the risky asset is assumed to be proportional to standard deviation of the asset to the power of a parameter "a." Also, the forecast of the standard deviation denoted by sigma-prime is assumed to be correlated with the true standard deviation and is denoted by "b." The model's results are that under a stable risk policy the weight in the risky asset is equal to the target standard deviation divided by the forecasted standard deviation, sigma-prime. Under the optimal strategy the weight in the risky asset is a constant divided by sigma-prime to the power 2-a. Furthermore, the model shows that the Sharpe ratios for the optimal and the constant risk allocation strategies are higher than the Sharpe ratio for the static policy strategy. Thus, the main finding is that the static policy allocation is a third-best strategy. The stable risk allocation strategy is a second-best strategy but it presents the advantage over the optimal strategy that it does not require any knowledge of the risk-return relationship.

6 Mark Kritzman

Mark Kritzman and Sebastien Page in their joint presentation address the problem of determining optimal portfolio weights when asset prices are dynamic. Institutional investors usually employ mean–variance analysis to determine portfolio weights, but almost immediately upon implementation these weights become obsolete, as asset prices have changed. In an ideal world, institutional investors would simply rebalance their portfolio weights continuously. However, with the presence of transaction costs, rebalancing portfolios continuously to restore the optimal weights may prove to be too costly of a strategy. Most institutional investors address this problem by rebalancing the portfolio weights at fixed periodic intervals or when the size of the misallocation is sufficiently large. An accurate way of solving the rebalancing problem is to implement a dynamic programming solution. Investors want to minimize the recursive cost function of the loss in utility from having suboptimal weights, the transaction costs incurred by changing those weights and of the future costs. This straightforward approach faces strenuous computational complications. For example, to rebalance a portfolio among three assets in increments of 1%, the number of calculations one would need to perform in order to solve the dynamic programming problem for a one-year horizon with 12 time steps (monthly monitoring) approaches 15 billion. For ten assets the computations jump to 10 octillion (short-scale) calculations. An alternative approach is devised by using the Markowitz and van Dijk (2004) heuristic. This method consists in rewriting the cost function as the loss in utility from having suboptimal weights, the transaction costs incurred by changing those weights, as before, and a new term which proxies for the future costs and is proportional to the squared deviations of the current portfolio weights from the optimal weights. This method allows the dynamic programming problems to be circumvented with even hundreds of assets. Testing the method on a portfolio composed of domestic equities, domestic fixed income, foreign developed equity, foreign bonds and foreign emerging equity, it is found that the Markowitz–van Dijk heuristic works as well as the dynamic programming method. Because the Markowitz–van Dijk heuristic has the same overall efficacy and is scalable to potentially hundreds of assets it opens the door for several new applications. For instance, managers could use this new method to optimize the tradeoff between tracking error and transaction costs and quantitative asset managers could use it to minimize alpha decay between rebalancing dates. Furthermore, the Markowitz–van Dijk heuristic could prove beneficial to plan sponsors since they are continuously confronted with asset mix rebalancing decisions.

7 Thomas S.Y. Ho

Thomas Ho proposes a new volatility risk measure called the key rate vega which is an extension of the vega measure to the buckets along an implied volatility function. To manage the risk of an interest rate contingent claim practitioners need both duration and vega measures, which are the instruments' sensitivities to the shift in the swap curve and the volatility surface, respectively. To manage interest rate risks practitioners use the duration buckets along the yield curve called key rate durations. However to date, there are significant computational challenges in determining the vega buckets for interest rate derivatives. Thomas Ho in his presentation shows how key rate vegas can be calculated using the result from Ho and Mudavanhu (2007) that any interest rate contingent claim can be valued by the swap curve and the implied volatility curve. Key rate vegas are then simulated for a set of at-the-money swaptions. The results suggests that when the volatility surface is stochastic, delta hedging of some interest rate contingent claims with vanilla swaps may not be effective because of the vega effect. Instead the study shows that both swaps and swaptions should be used even for dynamic hedging. Furthermore, in hedging the volatility risk, one cannot simply use one vega measure. To manage the volatility risk appropriately, one needs to measure the value sensitivity of an option to the change in the implied volatility function. The change may be specified at the key rate points on the yield curve.

8 Lisa Goldberg

In her presentation, Lisa Goldberg develops a risk calculator that forecasts the portfolio loss surface using high frequency data. Traditional risk measures such as Value at Risk (VaR) and Expected Shortfall (the expected loss given that VaR is breached) use a fixed horizon to forecast loss. The loss surface, however, is the distribution of portfolio loss at all future horizons, which requires a model of the dynamics of short-horizon returns. The loss surface allows for a common perspective on different aspects of portfolio risk since the more traditional risk measures such as volatility, VaR, and Expected Shortfall can all be recovered from it. The risk calculator uses the fact that loss over a longer horizon can be expressed as a sum of shorter horizon losses with the added complication that the latter are not independent and identically distributed. Thus, the strategy to forecasting the loss surface is to first collect a history of daily and to extract from this raw data an independent and identically distributed series of loss innovations by correcting for serial autocorrelation and heteroscedasticity using a one-lag autoregressive model with time-varying volatility. The distribution of one-day loss innovations is then estimated using a semi-parametric model called "peaks over thresholds" and a Fourier transform is used to compute its characteristic function. The properties of the autoregressive model are then used to express the T-day delta loss in terms of the one-day loss innovations. The T-day delta loss is simply the volatility-scaled difference between the portfolio loss over T-days and the autocorrelation coefficient scaled one-day loss innovations. The characteristic function of the T-day delta loss is then computed by taking the product of the characteristic function of one-day loss innovations. The next step is to recover the T-day delta loss distribution from its characteristic function using the inverse Fourier transform and to restore the effects of serial correlation and heteroscedasticity to obtain the distribution of T-day loss, which is the loss surface. The risk metrics obtained from this loss surface, labeled EVT for Extreme Value Theory, are tested empirically against other benchmarks such as the VaR and ES metrics obtained using the RiskMetrics (RM) methodology and other factor models. Using 76 value-weighted portfolios over the period 2000–2007 the 95% and 99% VaR and ES out-of-sample forecasts are calculated for a one-day, five-day, and ten-day horizon. At the one-day horizon, RM performed best for 95% VaR with EVT falling a close second and the other factor models falling far behind. For the one-day 95% expected shortfall forecasts, EVT dominated all other models. At the 99% level EVT and RM performed equally well, however the EVT model was much less skewed toward under-forecasting than the alternatives. The results are similar if one examines the five-day or ten-day horizon. Overall, these empirical results make a compelling case for the application and further development of the EVT approach in forecasting portfolio losses.

9 Jan Loey's

Jan Loey discusses the current turmoil affecting the financial system by showing the potential sources of the crisis and some solutions to it. The initial shock to the financial system started in June, 2007 when two Bear Stearns hedge funds that had invested in US subprime mortgages reported substantial losses. These losses induced the two major credit rating agencies to downgrade a large number of funds with similar holdings to the Bear Stearns funds, which severely damaged investor confidence. There are five potential sources to the current crisis. First, one source of the current crisis is a technical correction. Equity and other risky markets usually pull back for short periods of time lasting from one to two months. Because market corrections are usually short lived it is unlikely that the current crisis is a correction. Second, one source of the current crisis is the US housing market crash which has been the main source of the subprime losses and the ratings crisis. However, US residential construction had been in a period of contraction for over a year prior to the current crisis and if it were the main source of the crisis the damage should have occurred earlier. Third, another source of the crisis is a lack of financial liquidity. In this view, the sudden fall in value of the subprime mortgages made investors temporarily uncertain about the true value of the asset-back securities, which created a liquidity crisis. Under this view the central banks can easily remedy to the situation by providing liquidity and rate cuts. Fourth, one view is that the cause of the current crisis lies in the high leverage rates of US households, and the growth in bank off-balance sheet securitization such as CDOs and CLOs. Fifth, according to Jan Loey a more convincing argument is that there is currently a rating crisis causing the financial flows from savers to borrowers to be disrupted. One of the largest financial flows from savers to borrowers is in asset-backed securities that have been financing half of consumer net new borrowing. These asset-backed securities (ABS) rely on credit ratings as their main gauge of value and the sudden realization that investment grade ABS could face significant losses forced many investors to exit that market.

Given that the main source of the current crisis has been established as a ratings problem. Jan Loey discusses the benefits of potential solutions. One solution is for the Federal Reserve Bank to cut interest rates but this will not solve the underlying informational problems that under-resourced investors are faced with in trying to assess the value of ABS securities. The second solution would then be to reform the credit ratings system on ABS. Although the credit ratings system on ABS needs reform this solution is not likely to have immediate consequences. Another potential solution is the gradual investment in resources needed for investors to set up their own value-discovery system to invest again in ABS and the use of government sponsored enterprises as potential investors in ABS. Finally, over time banks will want to reestablish balance sheet discipline by removing off-balance sheet items. Since none of these solutions are short-term fixes the resolution of the current financial crisis is likely to take years rather than weeks or months.