JOIM Conference Series
San Francisco, Spring 2007

Myron Scholes Keynote Speaker

Myron Scholes in his keynote address examines the issue of pension fund investing for defined benefit retirement plans. In particular, Myron Scholes asks the question of what is the equilibrium for pension funding and the repercussions that changes in accounting rules could have on the economy as a whole. Myron Scholes first discusses the taxonomy of different categories of return generating mechanisms. First, returns can be generated through systematic exposures or Beta which is the strategy employed by index funds or ETFs. Slightly related are some funds that seek to generate returns through forecasting factors such as the market, currencies, or other macroeconomic factor. Second, some funds seek to generate returns through Alpha by essentially going long undervalued securities and shorting overvalued securities relative to a discounted cash flow model. Third, Myron Scholes introduces Omega where returns are generated from providing liquidity. The fact that liquidity providers generate returns can be seen in the real estate market with active or quiet markets. An illustration of how liquidity matters would be a person selling a house on Ebay at 2:00am with a half-hour window. The price that house will fetch is likely to be inferior to the intrinsic price. Since the price of liquidity can fluctuate some funds can make money by providing liquidity when it is most valued thereby generating Omega returns. Related to liquidity provision but slightly different is risk-transfer. Liquidity is usually thought of as buying an asset and having to liquidate it at a discount whereas risk-transfer has the direction of hedgers that are willing to pay speculators (providers of liquidity). A model of valuation of liquidity is needed for funds to better identify these omega opportunities. Furthermore, investing in Omega funds can offer returns that have little correlation with major indices and hedge funds.

Having set the stage for investments that offer little correlation with major indices, Myron Scholes discussed the relevance of Omega (risk-transfer) in defined benefit retirement. Finance theory developed by Irwin Tepper and Fischer Black tells us that pension funds should invest in bonds to enhance corporate tax benefits. Yet, pension funds have invested significantly in equities with the belief that in the long-run equities outperform bonds and the pensions' liabilities are long dated. However, substantial losses in the equity markets in 2000 left many defined benefit plans under-funded. This has led several European countries but notably the United Kingdom to pass regulation that requires corporations to report under-funded benefit plans on their balance sheet. Thus, there has been a great impetus by corporations to reduce their deficits and buy long-term bonds to match their assets with their liabilities. This has caused the price of long-term bonds to increase and yields to fall, as the supply of these bonds was scarce. In turn, this has generated a movement by corporations to issue long-term debt, as it has been cheaper to do so.

Kenneth Singleton

Kenneth Singleton shows how the term-structure of sovereign debt CDS spreads can be used to extract both the probability that credit event (such as default) occurs and the loss given that credit event actually occurred. A CDS contract provides insurance against

credit events where the buyer of protection pays the default swap premium to the seller of protection. If a specified credit event is triggered, then the protection seller covers the losses incurred by the insured. The sovereign CDS market has gained substantial breath and liquidity in the last five years although it remains small when compared with the corporate CDS market. For instance, the top five sovereign names account for 40% of all quotes, while the top five corporate names account for less than 8% of quotes. However, on the bright side, the sovereign CDS market has led its corporate counterpart in that a broad range of maturities are actively traded. This makes the full term structure of CDS spreads available for inferring default and recovery information from market data. Using data from March 2001 to August 2006 on the CDS spreads of Mexico, Turkey and Korea, three countries with different geopolitical and credit rating characteristics the time-series properties of the risk-neutral mean arrival rate of a credit event is explored. First, Maximum-Likelihood estimation is used to identify the mean arrival of a credit event holding the loss-given-default constant at 0.75 which is the market practice. Second, the same estimation technique is used but the loss-given-default is also a parameter to be estimated. In the latter case, the likelihood functions call for much smaller values of loss-given-default for Mexico and Turkey (of around 0.25) but similar to the market convention of 0.75 for Korea. A one-factor model whereby the mean arrival rate of a credit event is assumed to follow a lognormal process is shown to capture most of the variation in the term structures of spreads. The results suggest that, for all three countries, there are systematic, priced risks associated with unpredictable future variation in the default intensity. Moreover, the time-series of the effects of risk premiums on CDS spreads co-vary strongly across countries. Regression analysis shows that the VIX option volatility index (proxy for event risk), the spread between the ten-year return on U.S. BBB-rated industrial corporate bonds and the six-month U.S. Treasury bill rate (viewed as a measure of both U.S. macroeconomic and global financial market developments), and the volatility in the own-currency options market, are important determinants of the risk premiums. This evidence is consistent with premiums for credit risk in sovereign market being influenced by spillovers of real economic growth in the U.S. to economic growth in other regions of the world. Equally notable are the findings that during some sub-periods, a substantial portion of the co-movement among the term structures of sovereign spreads across countries was induced by changes in investors' appetites for credit exposure at a global level rather than to reassessments of the fundamental strengths of these specific sovereign economies.

Darrell Duffie

Darrell Duffie, in his presentation shows how the intra-day allocation and pricing of overnight loans of federal funds reflect the over-the-counter (OTC) nature of that market. While there is a significant body of research on the microstructure of specialist and limit-order-book markets, most OTC markets do not have comprehensive transactions-level data available for analysis with the federal funds market being a rare exception.

Specifically, the study shows the determinants of the likelihood that some bank borrows from some other bank during a particular minute of a business day transactions-level data from Fedwire on the top 100 commercial banks by transaction volume on the business days of 2005. The dataset consisting of over 225 million transactions permits the

construction of real-time balances for each institution allowing the sender and receiver of both payments and loans to be identified. A federal funds transaction consist in an electronic request by a financial institution to the Federal Reserve Banks via its Fedwire Funds Service to debit by a certain amount its federal funds account in favor of another financial institution. Federal funds loans are not collateralized therefore exposing the lending institution to the risk that the borrowing institution defaults. This last feature might explain in part the OTC nature of the federal funds market as opposed to a centralized market that requires more homogenous assets.

In modeling the allocation and pricing of overnight loans of federal funds it is important to understand the imperatives of market participants. At the end of each business day each financial institution is required to hold a positive balance or face the consequence of hefty overdraft fees. However, the Fed does not provide interest on the balance, therefore, financial institutions also have an incentive to maintain as little balance as possible. This last point is reflected in the fact that the total amount of reserves held by financial institutions is roughly $17.3 billion whereas the total daily amount sent on Fedwire is over $2.3 trillion per business day.

The study, first, examines the likelihood that some bank borrows from some other bank during a particular minute of a business day. The finding is that this likelihood depends on the prior trading relationship between the given two banks, the extent to which their balances at the beginning of the given minute are above or below their normal respective balances for that time of day, the overall levels of trading activities, and the amount of time left until the end-of-day balances are monitored for reserve-requirement purposes. The study then examines how the interest rate is negotiated in a particular trade. The finding is that the interest rate depends on some of the previous incentive-related variables but also on a measure of the credit quality of the borrowing financial institution.

Paul Kupiec

Paul Kupiec presents a new generalized single common factor model of portfolio credit risk that features both a stochastic exposure of default and loss given default rate. Generalized single factor models of portfolio credit losses, such as the one of Vasicek (1987), provide an approximation for the loss rate distribution in which the dependence among individual defaults is driven by a single common latent factor. The model assumes that the unconditional probability of default is fixed and known and that the exposure at default and the loss given default are non-stochastic. These single factor models have been widely applied in the financial industry and is the model that underlies the Basel II Advanced Internal Ratings-based approach for setting banks' minimum regularly capital requirements. However, the model's assumptions that the loss-given-default and/or the exposure-to-default are non-stochastic have been the subject of recent criticisms. Indeed, empirical studies show that the loss-given-default (LGD) and the exposure-at-default (EAD) are higher in periods when default rates are elevated. The presentation thus shows how the model can be extended to incorporate obligors with EADs and LGDs that are correlated random variables. In this extension, default is a random event driven by a compound latent factor as in the standard model but two additional compound latent factors are introduced to drive correlations among individual creditors' EADs and LGDs

using the inverse integral transformation. A closed-form expression for the inverse of the portfolio's conditional credit loss distribution is derived for cases when LGD and EAD probability density functions are continuous. However, when these distributions are made to be discrete, the integral transformation is inadmissible, and an alternative approach must be used to derive the asymptotic portfolio's loss distribution. The alternative approach consists in constructing the loss rate distribution using a step function to approximate the underlying LGD and EAD distributions, which yields a closed-form solution. Using the step-function representation, the calculations needed to approximate the portfolio loss rate distribution are straightforward and can easily be programmed in a financial spreadsheet. The model's portfolio loss rate distributions are illustrated for representative wholesale and retail portfolios. The results show that the additional systematic risk created by positive correlation among individual account EAD and LGD realizations increases the skewness of an asymptotic portfolio's loss rate distribution. This increase in skewness increases the measured risk in lower tranches of collateralized debt obligations (CDOs) and shows the need for larger economic capital allocations than those calculated with the Vasicek model.

Peter Carr

Peter Carr explores how unique arbitrage-free prices of set of default-contingent claims can be extracted using the entire CDS curve as given along with the term structure of interest rates. The default-contingent claims that are the focus of the study include any claim promising to make fixed payments over time according to a pre-set schedule such as a defaultable annuity. Although the model attempts to make as few probabilistic restrictions as possible it turns out that deterministic interest rates and constant recovery paper are needed in order to extract the initial value of the premium leg of a CDS. Of particular interest to institutions are the positions that are needed in order to replicate the payoffs of the default-contingent claims. The hedging instruments that are used to replicate those target claims consist of a cash account and CDS contracts of all maturities. The positions in the CDS are static and to hedge the payoff of a target claim the hedger needs to initially assume a position in zero cost CDS of all maturities and next he needs to set up a bank account with a balance that matches the theoretical value of the target claim. As time evolves, CDS values deviate from zero, and as a result, the bank account balance deviates from the theoretical value of the target claim. Prior to any default by the reference entity, the positions in the CDS and in the target claim require cash inflows or outflows, which are financed out of the bank account. Since the hedge is static in nature the initial CDS positions are kept constant throughout cutting down on transaction and monitoring costs. The replication portfolio weights are determined via dynamic programming, numerically solving backward ordinary differential equation for the survival contingent bank balance. While deterministic interest rates and a constant recovery rate on the CDS is assumed this static hedging strategy is otherwise robust in that no further assumptions on the process triggering default are needed. The results are thus consistent with both reduced-form and structural models of default.

Mark Seasholes

In his presentation, Mark Seasholes investigates the link between market-maker inventories and asset prices using a unique dataset of NYSE specialists inventories. Theoretically, liquidity suppliers and market-makers provide immediacy to less patient investors but have limited inventory carrying and risk bearing capacity. Market-makers accommodate trades when there is buying pressure by holding less than optimal levels of inventories and when there is selling pressure by holding more than optimal levels of inventories. In turn, market-makers get compensated for this service by buying (selling) the assets at a discount (premium) relative to future prices. Thus, empirically, one should observe that when inventories are large, liquidity suppliers have taken on risk and future prices should increase. Similarly, when inventories are low future prices should decrease. Empirical studies of the link between inventory and prices has typically taken place on a intraday horizon or on an interday horizon using a limited number of securities for relatively short periods. In contrast, this study uses daily inventory data for an 11-year period (from 1994 through 2004) on each stock on the NYSE using internal data from the specialist summary file (SPETS). The key findings are that consistent with theory the specialists' positions are negatively correlated with contemporaneous price changes (correlation of -0.71) and positively correlated with subsequent changes. If one were to create a portfolio that is long on stocks with the highest inventory positions and short on stocks with the lowest inventory positions the associated average returns are 0.10% the next day and 0.33% for the next 5-days. While these returns seem large, it is important to note that specialists do not disclose their inventory positions, so predictability based on inventories comes from non-public information. Furthermore, using Fama-MacBeth regressions the study finds that past inventories along with past returns and past order imbalances are able to predict return reversals.

Francis Longstaff

Francis Longstaff shows how information in collateralized debt obligation (CDO) prices can be used to study market expectations about how corporate defaults cluster through time or are correlated across firms. CDOs are financial claims to the cash flows generated by a portfolio of debt securities (or equivalently, a basket of Credit Default Swap (CDS) contracts. CDOs can also be viewed as being the credit-market counterparts to the familiar Collateralized Mortgage Obligation (CMO) structure. Since their introduction in the mid-1990s, the market for CDOs has grown dramatically and stands currently at an excess of $2 trillion with issuance in 2006 nearly doubling. The important drivers of growth have been the creation of standardized CDX and ITraxx indexes as well as the contemporaneous growth in the credit derivatives market as a whole. The study uses a three-factor portfolio credit model motivated by prior evidence that credit spreads are driven by multiple factors. Furthermore, portfolio losses are triggered by the realizations of three independent Poisson processes each with their own intensity and jump size. The jump sizes are 0.4%, 6% and 35% respectively. The first jump can be interpreted as representing the idiosyncratic default of a single firm, the second jump could be interpreted as the joint default of firms in a sector or industry, and the third jump would be a catastrophic economy-wide event. This model is then taken to the data using a unique dataset of CDX index tranches for the 2003-2005 period. The main findings are

that to properly model CDO prices all three factors are needed with on average, 65 percent of the CDX spread due to firm-specific default risk, 27 percent to clustered industry or sector default risk, and 8 percent to catastrophic or systemic default risk. However, firm-specific default risk has begun to play a larger role recently. The three-factor model is a very good fit to the data with root mean square errors less than one basis point. Thus, the study provides direct evidence of market expectations of default clustering.

Brett Trueman

Brett Trueman in his presentation asks the following research question: do past stock market winners exhibit a predictable return pattern around their earnings announcements. The research question is motivated by recent findings by Trueman, Wong and Zhang that examine returns around internet stocks' earnings announcements during the 1998 to 2000 period. In that study the finding was that large average abnormal returns preceded earnings announcements and were followed by a return reversal subsequently. Since that sample period corresponded to a time when internet stocks were rising rapidly the research begs the question as to whether there is a general phenomenon whereby winning stocks exhibit large unexplained increases prior to earnings announcements with subsequent reversals. To answer the question, a sample of firms consisting of all quarterly earnings announcements on available on COMPUSTAT between 1971 and September 30, 2005 is collected with the constraint that firms are listed on CRSP, have a December 31st fiscal year-end and whose stock price at the end of the previous quarter is at least $5. The final sample consists in 293,630 firm-quarter observations. Using this sample firms are grouped into deciles based on the 12-month returns prior to the quarter preceding the earnings announcements. Once deciles are formed the 5-day pre-announcement and 5-day post-announcement market-adjusted returns are calculated. The main finding of the paper is the almost monotonic increase in pre-announcement average market-adjusted returns from lower to higher deciles. Moreover, the average market-adjusted return for the top decile, 0.83 percent, is more than 50 percent greater than that of the ninth decile and is almost three times as large as the average pre-announcement market-adjusted return of 0.3 percent over the entire sample. The paper also finds that if the top decile itself is divided into sub-deciles there is also a monotonic relationship. The average pre-announcement market-adjusted return for the top sub-decile, 1.36 percent, is over 60 percent higher than that of the top decile as a whole. The top sub-decile's average post-announcement market-adjusted return of -1.75 percent is over twice the size of that for the top decile. These results are robust to (a) the fact that COMPUSTAT earnings dates are not always accurate, to (b) fact that multiple announcements might be occurring on the same day, and (c) transaction costs. There are two potential explanations for this phenomenon. The first possibility is that in the few days before an earnings announcement analysts raise their earnings forecasts to levels that are unjustifiably high and that investors take these revisions at face value. Subsequently, when earnings are released and fall short of expectations, the price drops back down. This explanation depends on investors not learning over time that analysts' forecasts are biased upward prior to earnings announcements. Using IBES data the study finds that in less than 2% of the sample there is both an upward revision in analysts' forecasts prior to earnings announcements and a negative earnings surprise subsequently. Thus the results do not

support the first explanation. The second possible explanation is that limited time and resources preclude individual investors from considering all possible equity investments. Consequently, they are more likely to buy stocks that draw their attention. This explanation is tested by calculating the abnormal order imbalance for small, medium-sized, and large traders. Since smaller investors are arguably less sophisticated they are more likely to buy stocks with strong prior returns just before their earnings announcements. Consequently, we would expect to observe an unusually large number of buyer-initiated trades relative to seller-initiated trades in the pre-announcement period for these traders, but not necessarily for larger ones. Once earnings are released and the focus shifts from these stocks, this positive abnormal order imbalance should disappear.

The results are consistent with this second explanation. During the pre-announcement period small and medium-sized traders evidence a significantly positive abnormal order imbalance but an insignificant one for large traders. In the post-announcement period the positive abnormal order imbalances of the small and medium-sized traders disappear. This suggests naïve investors are at least partly responsible for the phenomenon.

Andrew Lo

Andrew Lo presents a new metric for measuring the performance of actively traded portfolios. The new measure is essentially an active/passive decomposition that takes into account the forecast power of the portfolio manager explicitly. Traditional measures such as Alpha, Beta, Volatility, Tracking Error, the Sharpe ratio, and the Information Ratio are popular are now standard metrics to gage the value-added of long portfolio managers. However, these measures have not gained much popularity among investors of absolute-return strategies that are provided by hedge-fund managers. One reason might be cultural in that the growth of the mutual fund industry largely coincided with that of portfolio theory and the benefits of diversification. Another reason is that traditional performance measures are static in nature and do not capture the dynamic and predictive nature of active investment strategies. For instance, Alpha, Beta, Tracking Error and the Information Ratio are all functions of parameters that are defined at one point in time yet the central focus of active investment strategies revolves around the relation between returns at multiple points in time. For example, the most admired portfolio managers are revered for their ability to foresee certain market trends well in advance of the public or detect mispriced securities and exploit them ahead of the market, or to enter or exit certain investments before others recognize the opportunities. In every case, these investment skills involve forecasts or predictions that the standard performance measures do not explicitly capture. The active/passive decomposition that Andrew Lo proposes consists in decomposing the expected return of a portfolio into two components: one that depends on the average values of the portfolio weights and asset returns and the other that depends on the correlation between the portfolio weights and returns. The former component measures the performance of a passive strategy, whereas the latter component measures directly the value of active management. The successful manager will generally have portfolio weights that are positively correlated with returns generating a positive contribution to the portfolio's expected return. Furthermore, the correlation between portfolio weights and returns is a direct measure of the manager's timing ability because

the weights reflect prior information. This Active/Passive decomposition is easy for managers to implement and does not even require position-level information. Only average portfolio weights and individual asset average returns are needed to perform the decomposition. In addition to being a performance measure the Active/Passive decomposition can yield additional insights if asset returns are assumed to follow a linear asset pricing model. In particular, a portfolio's expected return can be decomposed into three separate components: security selection, factor-timing ability, and risk premia. While the first two components result from active management, the last component is passive.

Andrew Lo Tutorial

Andrew Lo presented a tutorial on risk management. The tutorial included a presentation of why risk management is an important activity and on the origins and current practices of risk. Andrew Lo then introduced the three P's that are important to risk management which are Prices, Probabilities and Preferences. These three P's can constitute a unified framework for thinking about risk management.

First, to illustrate the importance of risk management even to managers that seek Alpha returns, Andrew Lo showed how a manager that has access to a fairly mediocre investment that has expected returns of 5% and a standard deviation of 75% and a risk management protocol that limits the potential losses to 10%. The resulting expected return on the investment is 22.7%. Thus, successful risk management can generate alpha returns if the costs are properly managed.

Risk management has its origins in the OTC derivatives trading desks where there were complex securities such as options and other derivatives put together in portfolios with dynamic non-linear payoffs where the traditional investment portfolio measures such as tracking error, and beta did not make sense. The measure that was developed to quantify risk was VaR or Value-at-Risk. Value-at-Risk is an estimate, of how much one can lose from holding a position over a set horizon given a predefined confidence interval (such as 1% or 5%). VaR can be estimated with daily, monthly, or quarterly frequencies, normal or fat-tailed distributions, and using various estimation procedures for the mean and volatility of the distribution. Several limitations of VaR include the fact that VaR is generic in the sense that it doesn't depend on any security or market or make use of any information from those markets. The measure is static in that there is no time-series element to it or forecasting and it is parametric in that it makes distributional assumptions such as normality, which do not always reflect the reality of returns. Furthermore, the estimates are poor because VaR focuses on outliers that are not often observed. Improved risk management involves marking-to-market rather than marking-to-model, stress testing, creating scenario analyses, understanding the capital adequacy requirements involved with leverage, and understanding the time-series properties of the investments.

Andrew Lo then discusses an integrated framework of total risk management, which he calls the three P's, which stand for Prices, Probabilities, and Preferences. The three P's originated from basic microeconomics where the product of the marginal utility of good

and the price of the good as to be equal across all goods. With financial products the expected marginal utility times the price has to be equalized across all investments. Thus risk management involves investing in such a way that the marginal benefit of each investment opportunity is equated across all investment opportunities. This is the complete and integrated framework for thinking about risk because it suggests that across all possible investments the manager needs to invest in each activity so as to equate the marginal benefits across all possible activities. Thus, risk management involves gathering prices that are determined by the market, estimating probabilities, and determining risk preferences. In other words, a manager must consider how much does it cost to hedge, the likelihood that hedging will be needed, and how much to hedge.

The three P's are apparent in multi-factor risk models where one might want to hedge against a given factor. Multi-factor risk models allow the variation of returns to be decomposed into idiosyncratic and systematic risk and the systematic risk to be decomposed further into factors such as the market, currency factors, or even hurricanes and earthquakes. The manager then needs to determine how much exposure to a given factor is needed which is done according to her preferences.

Andrew Lo then discusses the need for new analytics in risk management that are needed because of the dynamic nature of active management by hedge funds and regime switching models. Regime switching models are models that are characterized by discrete "states" of the world each having its own statistical properties. A state can be a recession or an economic expansion or even different fiscal or monetary policies caused by political changes. Regime switching models can help to refine risk management models.

Finally, Andrew Lo showed how preferences could be modeled experimentally by asking managers what decision they would make given the choice of different gambles.