

Improving the Business of Risky Business

An *Ex Post* Evaluation of VaR's Statistical Input Assumptions

Michael Douglas Dechert and John Arthur Martin

Professor Emma Rasiel, Faculty Advisor

*Honors Thesis submitted in partial fulfillment of the requirements for Graduation with
Distinction in Economics in Trinity College of Duke University*

Duke University
Durham, North Carolina
April 19, 2010

Acknowledgements

We would like to thank Professor Emma Rasiel for all of her attentive and insightful advice throughout both the research and writing processes. Without the tireless help of Dr. Rasiel, this paper would not have been possible. We would also like to acknowledge the help of Ken deRegt, Chief Risk Officer, Morgan Stanley; Eduardo Canabarro, Managing Director, Market Risk Management, Morgan Stanley; and Edward Talisse, Executive Director, Operational Risk Management, Morgan Stanley, for their keen insight regarding VaR's practical application in modern-day risk management. In addition, we would like to thank Bill McMahon, Chief Risk Officer, Goldman Sachs, for his thoughtful perspective on the VaR model and the most recent financial crisis. Furthermore, we would like to thank Dr. Jonathan Wiener, former President of the Society for Risk Analysis, Duke University Professor of Law, Professor of Environmental Policy, and Professor of Public Policy Studies, for his opinions on a variety of current risk issues and on the importance of extreme tail events. Finally, we would like to thank both Dr. Lori Leachman, Professor of the Practice of Economics at Duke University, for her guidance on financial theory and academic writing and Aino Levonmaa, Visiting Assistant Professor of the Practice of Economics at Duke University, for her thoughtful explanations of alternative risk management tools.

Abstract

We evaluate various Value-at-Risk (“VaR”) statistical input schemes, comparing the relative effectiveness of dynamic volatilities and correlations to static assumptions based on historical data. Using existing data from a Trading Game in three market environments, we calculate traders’ portfolio VaR, *ex post*, under differing statistical input assumptions and then adjust their positions day-by-day to match their portfolio VaR and asset allocation under the original static assumptions. We find that the dynamic assumptions result in higher returns with lower volatility, as they compel traders to “come closer to home” during volatile environments, while still allowing them freedom to exploit favorable markets.

I. **Introduction & Literature Review**

The Value-at-Risk (“VaR”) framework is a risk management model that estimates traders’ risk exposure based on volatility and correlation assumptions. This mathematical model currently is by far the most widely used risk model at financial firms (Nocera, 2009; Berkowitz and O’Brian, 2002). To its proponents, its greatest appeal is that it provides a single, dollar-denominated figure, steeped in established statistical theory, that encapsulates risk exposure for a given portfolio. However, the effectiveness of the VaR framework is dependent upon the statistical input assumptions used. Without accurate estimates of current volatility and correlation, the VaR model will not appropriately limit traders’ positions or provide a useful depiction of their portfolios’ risk exposure. In this paper, we evaluate alternative VaR statistical input schemes, specifically comparing the relative effectiveness of rolling, dynamic inputs to long-term historical, static assumptions in three different market environments.

The autumns of 2006, 2007, and 2008 witnessed three distinct economic climates, with equity markets rising, remaining relatively neutral, and plummeting, respectively. During each of these three market environments, Duke University undergraduates (selected based on their performance on a general “market knowledge” test) participated in a three-month long “Trading Game.” For this Game, each participant (“trader”) is given an initial hypothetical endowment of \$1,000,000, with which he or she is able to speculate in ten different financial assets.¹ There is a limit imposed on each trader’s daily portfolio VaR as well as on each daily asset VaR. These original VaR measurements were estimated using historical, static volatility and correlation assumptions, referred to as “Base Case” assumptions in this paper. So, each day, the traders are restricted by eleven different limits –

¹ There were eleven assets in 2008 due to the addition of VIX, which was only traded lightly.

one for each of the ten assets as well as one overall portfolio risk constraint.² With the grand prize of a trip to Morgan Stanley’s equity and fixed income trading floors and the opportunity to meet with senior Morgan Stanley traders, the students had incentive to treat the game seriously.

The day-to-day trading data available from the Game in 2006, 2007, and 2008 present a unique opportunity to evaluate the success of various VaR statistical input schemes in these three different market environments (bull, neutral, and bear). Using this data, we carry out a “Horse Race” to determine the effectiveness of two dynamic volatility and correlation estimate schemes, relative to the Base Case assumptions. Subject to the dynamic VaR inputs, we modify each trader’s daily positions, *ex post*, in such a manner as to preserve their portfolio VaR and asset allocation by market value. We then aggregate the traders’ positions and compare the “Trading Floor”³ returns and the volatility of those returns in each market under each VaR input scheme. Finally, we evaluate each scheme’s resulting Sharpe ratio. Our hypothesis is that VaR inputs that are more dynamic in nature will better reflect changing market environments, and thus result in higher returns and lower volatility in any market conditions.

The Value-at-Risk model operates under a series of simplifying assumptions and is therefore limited in its application in the actual financial marketplace. First, VaR calculations assume a normal distribution for asset returns. However, extensive research has shown that financial markets demonstrate non-normal distributions (Richardson and Smith, 1993). For example markets are characterized by leptokurtic behavior, or the presence of “fat tails” – higher probabilities of extreme returns than a normal distribution would predict (Lucas,

² See Section 3 for specifics on how limits are defined.

³ For our purposes, the term “Trading Floor” refers to the combination of all individuals’ positions on a given day.

2000). Whereas a normality-based model predicts zero six-sigma S&P 500 returns since 1927, *forty-seven* actual such events have been recorded (Pine Cook Capital, 2008). More to this point, a left skewness has also been observed in asset return distributions (Sheikh and Qiao, 2009). This means that, in addition to having *fat* left tails (leptokurtosis), return distributions are also characterized by *long* left tails. This indicates a higher frequency of extreme, negative events relative to extreme, positive ones. Because fatter, longer left tails exist in the markets, the VaR model’s assumption of normality will inherently underestimate the probability of severe, negative swings in asset prices (Andersen *et al*, 2005).

Second, the VaR model provides no information about the scale of potential losses on days that fall outside of its predictive range (Beder, 1995). For instance, if a firm’s daily return is expected to lie within \$10MM 95% of the time⁴ (i.e. it has a daily portfolio VaR of \$10MM), \$10MM is simply the *minimum* amount the firm is expected to make or lose 5% of the time. VaR has no mechanism to estimate the magnitude of gains or losses that occur in this 5% range. Critics such as Taleb (2007, p.74) believe these models “to be wrong with infinite precision,” as they place importance on what will *probably* happen and ignore the very extreme situations—the so called black swans—that could *possibly* happen (Bogle, 2008). Though improbable, such events can leave a firm in financial ruin, and VaR cannot incorporate the scope of such events into its calculations.

Finally, a further deficiency of VaR has been exposed throughout the most recent financial crisis. Though assets vary considerably in their levels of liquidity, the VaR model fails to account for liquidity risk among assets. At the extreme, if there is *no* market price for an asset, then any risk model relying on volatility assumptions will not provide a meaningful

⁴ VaR can be constructed to report any level of confidence, and this 95% confidence level is merely an example of one possibility (here, using two standard deviations).

output. This can prove to be a critical deficiency of VaR for firms with heavy exposure to emerging markets, which are particularly prone to market failures (Bangia *et al*, 2001⁵). Even in developed markets, this can be an especially grave shortcoming of the model during times of tightened credit. Ethan Berman, CEO of RiskMetrics⁶, described VaR's lack of applicability in credit-constrained environments by describing it as essentially “a peacetime statistic” (Nocera, 2009).

Recently, much work has been done to account for this observed non-normality of returns with the intention of more accurately modeling expected distributions. A body of work known as Extreme Value Theory (“EVT”) provides a framework with which to consider the fat (and long) left tails discussed above (Gilli and Kellezi, 2006). EVT specifically considers the probability of observing rare, but high risk events, and a number of Extreme Value distributions have been suggested, such as the Generalized Pareto distribution. These Extreme Value distributions are used to fit the right and left tails, as distinct from the center of the distribution (Sheikh and Qiao, 2009). Additionally, recent research has focused on the observation that the relationship between the returns of multiple assets is not always linear, particularly at the extremes, as assumed by simple correlations. To account for this, copulas, functions that capture the interdependence of different assets' returns, can be applied to a joint distribution (Patton, 2006). Copula formulas provide a more robust and customizable means of modeling the joint evolution of multiple assets' returns than do simple correlations (Romano, 2002). Like Extreme Value distributions, copulas are adjustments that aim to more accurately estimate the distribution of portfolio returns, with

⁵ Bangia *et al* present an intriguing possible adjustment to VaR to help correct for this liquidity risk using the liquidity information embedded in bid-offer spreads.

⁶ RiskMetrics, hence spun off from J.P. Morgan, is the group widely credited with the creation of VaR.

particular emphasis on extreme events. Both Salvadori and De Michele (2007) and Yi and Bier (1998) have produced work employing the use of copulas in their respective fields.

The Conditional Value-at-Risk (CVaR) model provides a different approach to risk-measurement than does the traditional VaR model. CVaR measures the conditional expected loss of a portfolio, given the loss exceeds VaR (Krokhmal *et al*, 2001). For example, CVaR₉₅ is defined as the average real portfolio loss relative to the starting portfolio value in the worst 5% of Monte Carlo scenarios (Sheikh and Qiao, 2009). Unlike traditional VaR, CVaR quantifies the expected loss of a portfolio outside the predictive range of VaR using a weighted average of returns.⁷ This model can be applied to any assumed return distribution. However, when applied to distributions that assume fat left tails, CVaR will be higher than VaR, resulting in a tighter constraint. While these attempts to model non-normality and alter the VaR framework are helpful, the effectiveness of resulting models will still rely upon the statistical input assumptions used.

Despite its limitations and recent attempts to improve its structure, VaR is still ubiquitous in the financial world. Many large financial institutions are publicly traded, and consequently senior management is subject to boards, investors, and regulators who demand a depiction of the firm's current risk level (Jorian, 1996). VaR provides that "seductively" simple portrayal of risk that all of these parties can readily comprehend (Beder, 1995). Still, it is not possible to distill *all* possible risks that exist for financial firms into one easy picture. However, Bill McMahon, Chief Risk Officer, Goldman Sachs, noted that VaR is certainly a useful tool in the toolbox (Bill McMahon, personal communication, April 15, 2009). Perhaps, Richard Bookstaber, author of *Demon of Our Own Design*, put it best when he

⁷ In the above example, this weighted average would be calculated for the range of returns falling outside of the 5% threshold.

said, “If you put a gun to my head and asked me what my firm’s risk was, I would use VaR” (Nocera, 2009). With all of its imperfections, VaR is still the standard tool in modern risk management. As such, we see our research as adding to the understanding of how best to determine the statistical input assumptions for an important risk management instrument.

The structure of this paper is as follows. Section II provides a description of the data used in our research. Section III details the different Methodologies and describes the Horse Race and its calculations. Section IV presents the results of the Horse Race and analyzes the significance of our findings. Section V is the conclusion of the paper, placing our results in context and suggesting possible avenues of future research.

II. Data Description

The Trading Game was created by Professor Emma Rasiel as an extracurricular program for undergraduates at Duke University. Sponsored by Morgan Stanley, this Trading Game provided all of the data used in our research and calculations. Traders who participated in the Game during each of the three fall semesters (2006, 2007, and 2008) were permitted to speculate in the following ten assets: 2yr US Treasury Bonds, 10yr US Treasury Bonds, Baa Corporate Bonds^{8,9}, S&P500, Nasdaq, Dow Jones Industrial Average, Gold Futures, Oil Futures, \$/Euro, and \$/Yen. Traders in the 2008 Game were also permitted to trade VIX contracts. The Game allowed end-of-day trades only, and market prices and yields were recorded every evening using values from Bloomberg.¹⁰ Bond prices for 2yr US Treasuries, 10yr US Treasuries and Baa Corporate Bonds were calculated using constant maturity yields and an assumed coupon of 6%.

There were 79, 53, and 54 participants in 2008, 2007, and 2006, respectively. In 2008, the Game spanned the period from September 8 through December 5; in 2007, September 10 through November 29; in 2006, September 6 through November 29. On each day, the Game incorporated the assets' actual market returns in its calculation of each trader's Profit and Loss ("P&L"). Below, Tables 1 – 3 highlight each asset's starting, ending, minimum, and maximum values over the three Game periods. Charts 1 – 9 provide a graphical representation of equity, debt, and commodity performance over the time period of each Game. Additionally, Charts 10 – 12 graphically present rolling 90-day correlation

⁸ Via Moody's Seasoned Baa Corporate Bond Yield.

⁹ For the 2006 Game, 5yr US Treasury Bonds replaced Baa Corporate Bonds

¹⁰ All market data was taken from Bloomberg with the exception of Baa Corporate Bond yields, which were obtained from *Economagic*.

estimates¹¹ of the S&P500 with gold futures and with 10yr US Treasuries for each of the three Game periods.

In all three years, the Game was played under a detailed set of guidelines. Each trader began with a cash position of \$1,000,000 and was allowed to trade the Game's securities in discrete units only.¹² The traders were permitted to take on short positions; however, they were forced to pay a borrowing rate (benchmarked on the Federal Funds Target Rate) when their overall positions became net cash negative. During the 2008 Game period, this rate varied from 2.00% to 1.00%. During 2007, this rate varied from 5.25% to 4.50%. During 2006, this rate remained static at 5.25%. The traders were also subject to a static "transaction cost" of 0.04%. At the end of each trading day, this percentage was multiplied by the sum of the absolute values of the change in cash resulting from each asset traded, and the resulting dollar amount was subtracted from their end-of-day P&L. Finally, Game participants faced daily risk limits. A trader's daily VaR could never exceed \$50,000 for any individual asset, and his total daily portfolio VaR¹³ could not exceed \$100,000. All of these guidelines were designed with the intention of simulating a trading environment as realistically as possible.

¹¹ In Charts 10, 11, and 12, each point represents the correlation estimate between the two assets incorporating data over the time period $(t-90, t)$, where t is the date on which the point is graphed.

¹² It should be noted that, though participants in the Trading Game only traded security contracts in discrete units, we adjusted their trades, *ex post*, using infinitely divisible units in order to match exactly each trader's original daily portfolio VaR and asset allocation.

¹³ The calculation of daily portfolio VaR can be found in Section III.

Table 1: 2008 Market Statistics (9/8/2008 to 12/5/2008)

	2yr Yield	10yr Yield	Baa Yield	S&P500	Nasdaq	DJIA	Gold	Oil	\$/Euro	\$/Yen	VIX
STARTING	2.30%	3.67%	7.01%	\$1,267.79	\$2,269.76	\$11,510.74	\$806.90	\$101.46	\$1.41	\$0.0093	22.64
ENDING	0.92%	2.70%	8.71%	\$876.07	\$1,509.31	\$8,635.42	\$752.20	\$39.19	\$1.27	\$0.0108	59.93
MINIMUM	0.81% (12/04/2008)	2.55% (12/04/2008)	7.01% (09/08/2008)	\$752.44 (11/20/2008)	\$1,316.12 (11/20/2008)	\$7,552.29 (11/20/2008)	\$712.50 (11/12/2008)	\$39.19 (12/05/2008)	\$1.25 (11/12/2008)	\$0.0093 (09/08/2008)	22.64 (09/08/2008)
MAXIMUM	2.30% (09/08/2008)	4.08% (10/14/2008)	9.54% (10/31/2008)	\$1,267.79 (09/08/2008)	\$2,273.90 (09/19/2008)	\$11,510.74 (09/08/2008)	\$928.50 (09/29/2008)	\$103.86 (09/22/2008)	\$1.48 (09/22/2008)	\$0.0108 (12/04/2008)	80.86 (11/20/2008)

Table 2: 2007 Market Statistics (9/10/2007 to 11/29/2007)

	2yr Yield	10yr Yield	Baa Yield	S&P500	Nasdaq	DJIA	Gold	Oil	\$/Euro	\$/Yen
STARTING	3.87%	4.34%	6.47%	\$1,451.70	\$2,559.11	\$13,127.85	\$712.00	\$76.29	\$1.38	\$0.0088
ENDING	3.04%	3.94%	6.40%	\$1,469.72	\$2,668.13	\$13,311.73	\$798.80	\$91.18	\$1.48	\$0.0091
MINIMUM	2.99% (11/26/2007)	3.89% (11/26/2007)	6.23% (11/26/2007)	\$1,407.22 (11/26/2007)	\$2,540.99 (11/26/2007)	\$12,743.44 (11/26/2007)	\$712.00 (09/10/2007)	\$76.29 (09/10/2007)	\$1.38 (09/10/2007)	\$0.0085 (10/12/2007)
MAXIMUM	4.23% (10/12/2007)	4.70% (09/20/2007)	6.73% (09/20/2007)	\$1,565.15 (10/09/2007)	\$2,859.12 (10/31/2007)	\$14,164.53 (10/09/2007)	\$838.00 (11/08/2007)	\$95.81 (11/23/2007)	\$1.49 (11/26/2007)	\$0.0093 (11/26/2007)

Table 3: 2006 Market Statistics (9/6/2006 to 11/29/2006)

	2yr Yield	10yr Yield	Baa Yield	S&P500	Nasdaq	DJIA	Gold	Oil	\$/Euro	\$/Yen
STARTING	4.81%	4.80%	6.53%	\$1,300.26	\$2,167.84	\$11,406.20	\$648.20	\$72.66	\$1.28	\$0.0086
ENDING	4.63%	4.47%	6.14%	\$1,400.63	\$2,431.77	\$12,221.93	\$641.80	\$67.41	\$1.32	\$0.0086
MINIMUM	4.59% (10/04/2006)	4.47% (11/29/2006)	6.13% (11/28/2006)	\$1,294.02 (09/07/2006)	\$2,155.29 (09/07/2006)	\$11,331.44 (09/07/2006)	\$558.60 (09/15/2006)	\$63.46 (11/16/2006)	\$1.25 (10/13/2006)	\$0.0084 (10/10/2006)
MAXIMUM	4.91% (10/23/2006)	4.83% (10/23/2006)	6.53% (09/06/2006)	\$1,401.20 (11/17/2006)	\$2,460.26 (11/24/2006)	\$12,342.56 (11/17/2006)	\$648.20 (09/06/2006)	\$72.66 (09/06/2006)	\$1.32 (11/29/2006)	\$0.0086 (11/29/2006)

Chart 1: Fall 2008 Equity Market Performance

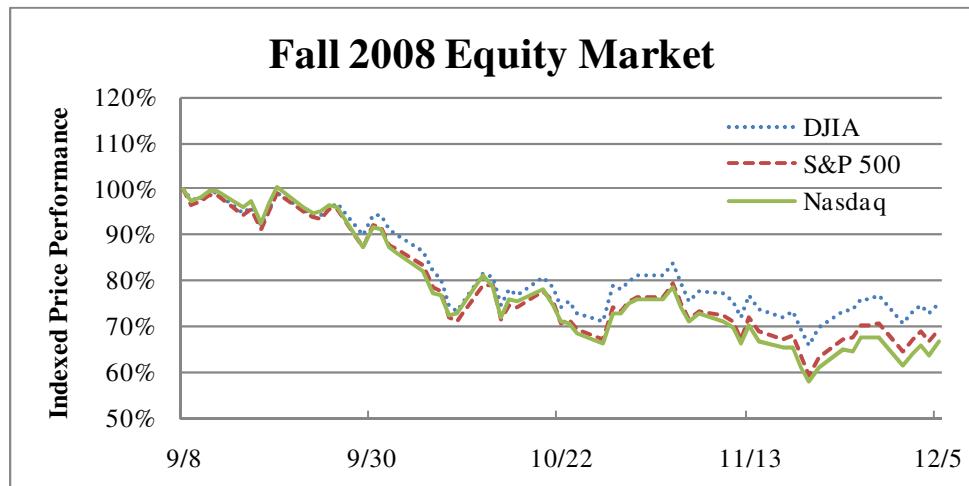


Chart 2: Fall 2007 Equity Market Performance

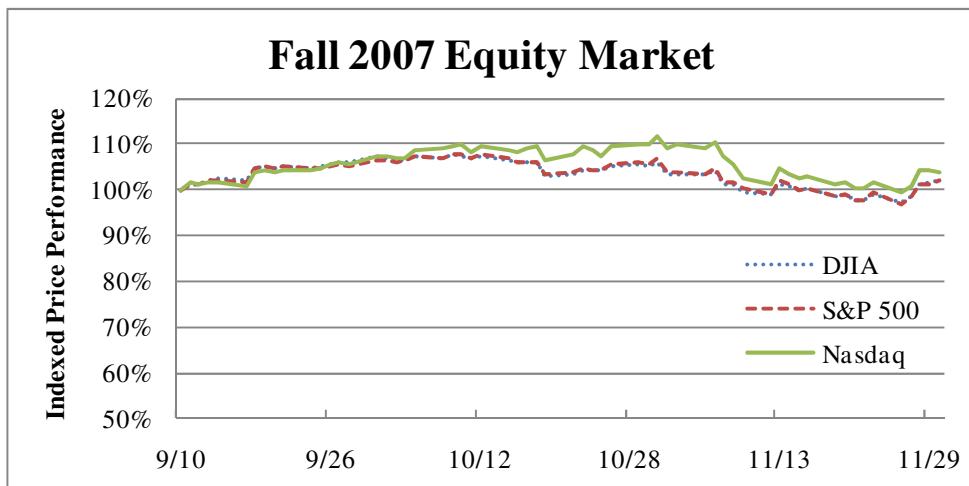


Chart 3: Fall 2006 Equity Market Performance

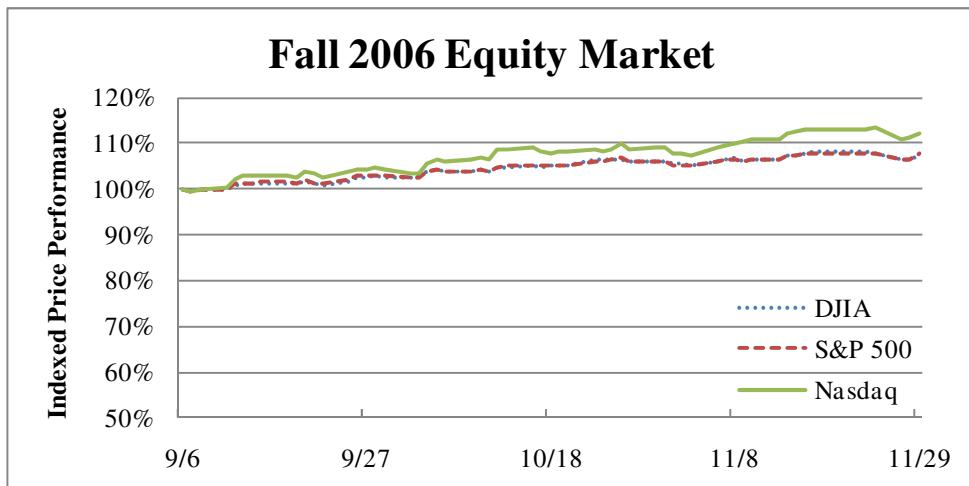


Chart 4: Fall 2008 Debt Market Performance

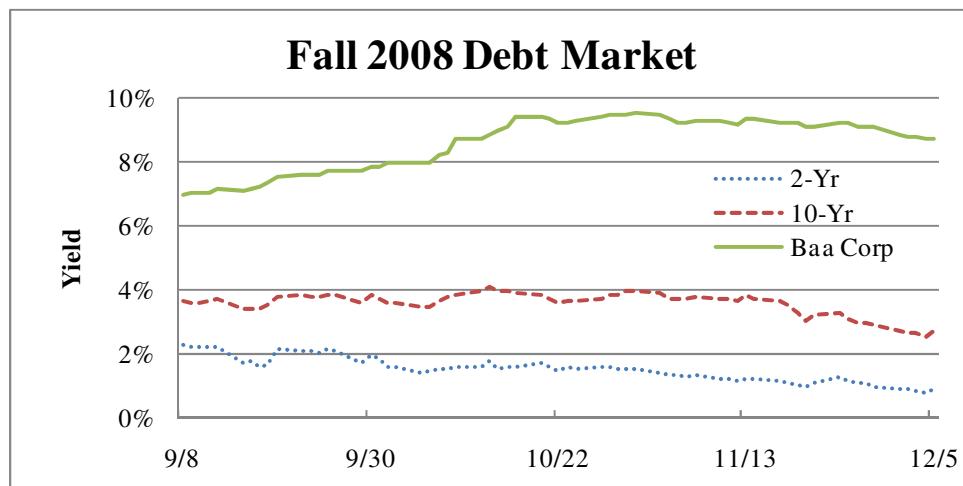


Chart 5: Fall 2007 Debt Market Performance

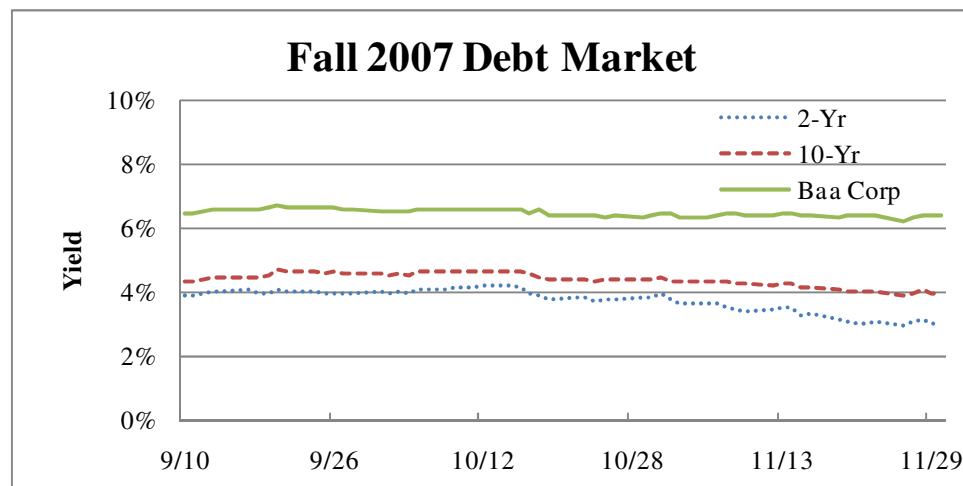


Chart 6: Fall 2006 Debt Market Performance

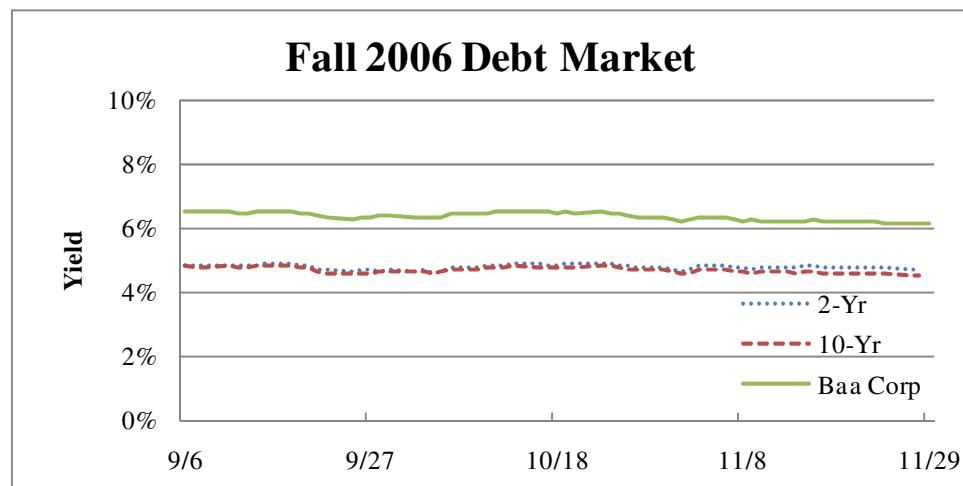


Chart 7: Fall 2008 Commodities Market Performance

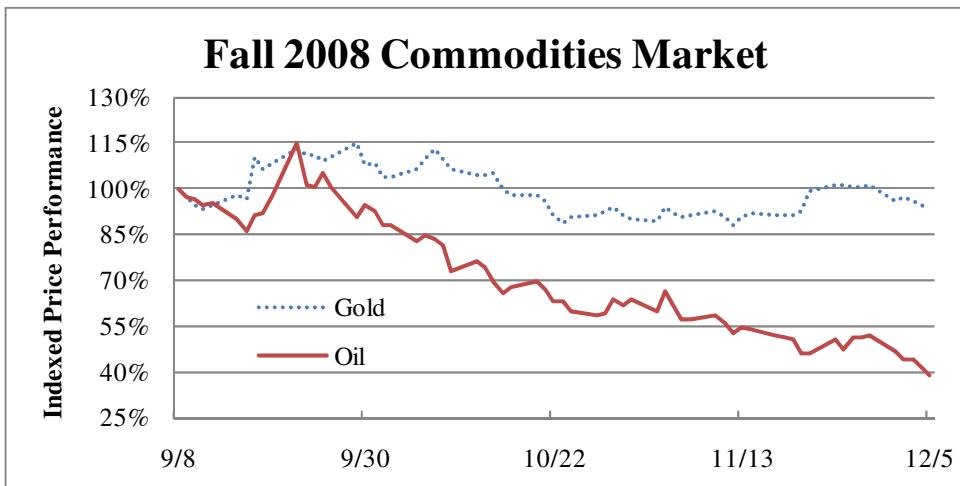


Chart 8: Fall 2007 Commodities Market Performance

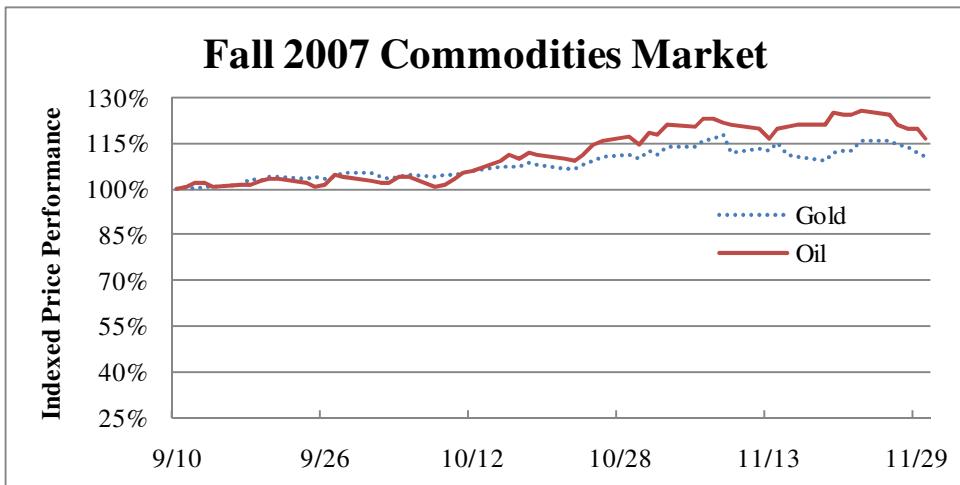


Chart 9: Fall 2006 Commodities Market Performance

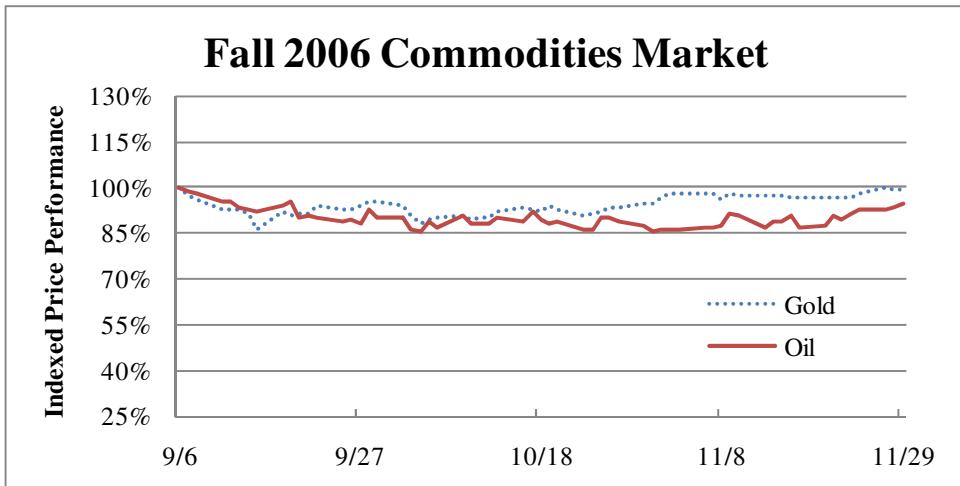


Chart 10: Fall 2008 Correlations

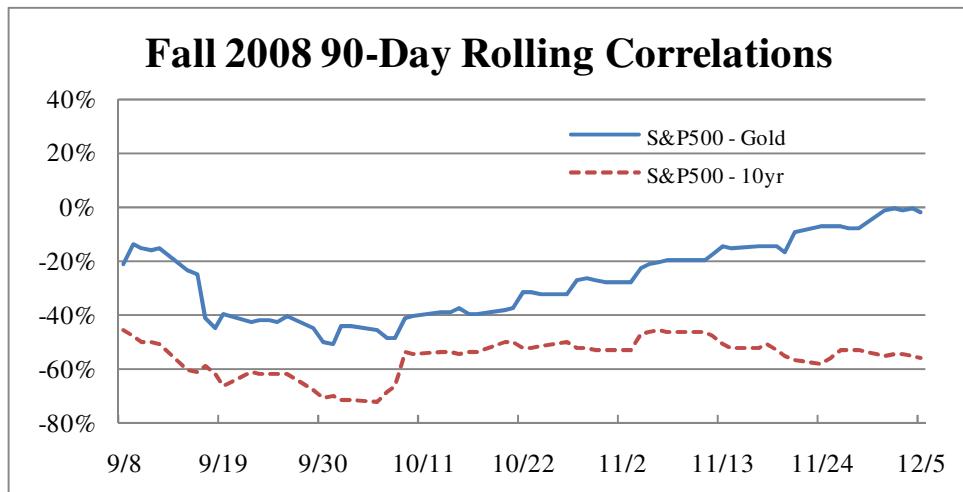


Chart 11: Fall 2007 Correlations

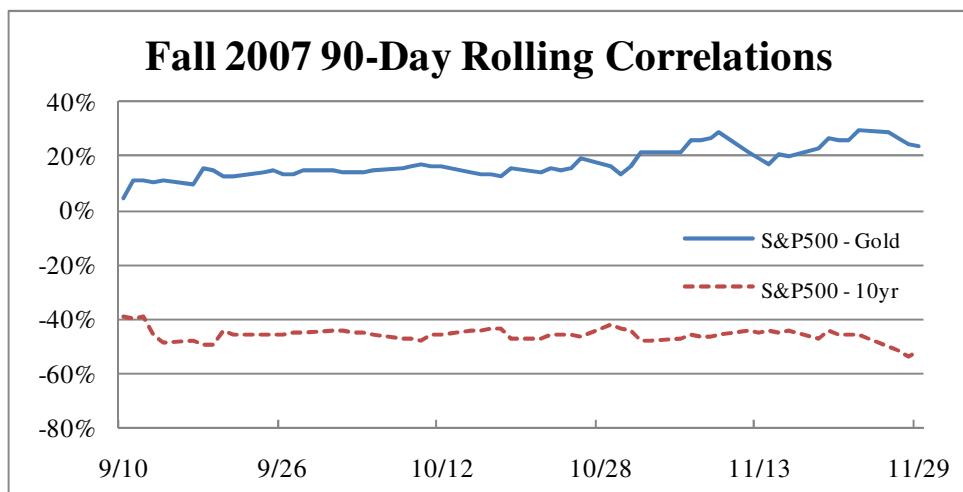
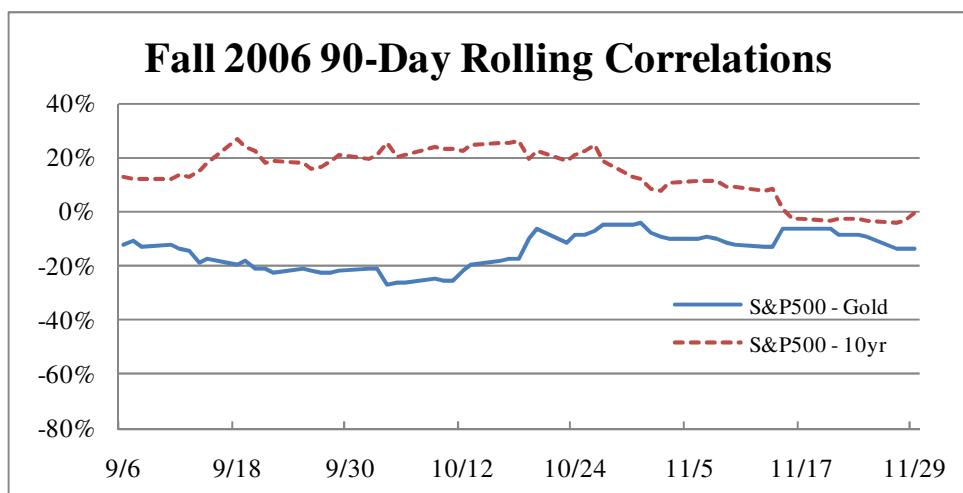


Chart 12: Fall 2006 Correlations



III. Methodology

Calculation of Daily Portfolio Value-at-Risk

The daily portfolio Value-at-Risk (VaR) represents the range in which a trader's daily P&L is expected to lie 68%¹⁴ of trading days (one standard deviation), while providing no information about how far outside that range his P&L will lie the other 32% of the days. To calculate this figure in a blended risky-riskless portfolio, one must first find the variance of the risky portfolio. The formula for this is:

$$(1) \quad \text{Risky Variance} = \sum_{j=1}^m \sum_{i=1}^n (w_i)(w_j)(\text{cov}(i,j)) ,$$

where w_i and w_j represent the proportion of the portfolio in asset i and j , respectively. Next, the weight of the risky portfolio within the entire portfolio is simply:

$$(2) \quad \text{Risky Weight} = 1 - \text{Riskless Weight} , \text{ where}$$

$$(3) \quad \text{Riskless Weight} = \text{Cash}/\text{Total Portfolio MV} .$$

Then,

$$(4) \quad \text{Portfolio Variance} = (\text{Risky Variance})(\text{Risky Weight})^2 ,$$

$$(5) \quad \text{Portfolio Daily Vol} = \sqrt{\text{Portfolio Variance}} , \text{ and}$$

$$(6) \quad \text{Portfolio Daily VaR} = \text{ABS}(\text{Total Portfolio MV} * \text{Portfolio Daily Vol})/\sqrt{252} .$$

¹⁴ This percentage can vary based on specific Trading Floor practices (i.e. 95% if two standard deviations are used).

General Model and Weighting Scheme

The general methodology of the experiment is to match each individual trader's original portfolio VaR levels and asset allocation on a day-by-day basis, regardless of which statistical input assumptions are used to measure portfolio VaR. We accomplish this by recording each trader's original daily portfolio VaR and then proportionally reducing or increasing each asset position equally to match that original VaR under different correlation and volatility assumptions. The weighting scheme calculations are as follows:

$$(7) \quad \text{Portfolio Daily VaR}_{new} = \text{Portfolio Daily VaR}_{old}, \text{ and}$$

$$(8) \quad w_{i,new} = w_{i,old},$$

for all i ; where i refers to the different assets (S&P500, gold, 10-year UST's, etc.), and w_i refers to the trader's portfolio weight in asset i . The constrained optimization system then finds k such that:

$$(9) \quad MV_{i,old} * k = MV_{i,new},$$

for all i , where k , the proportional reduction parameter across all assets, is a constant.

The inherent assumptions behind this approach are: (a) that traders revealed their total portfolio risk-tolerance by the maximum dollar amount they were willing to risk (i.e. their portfolio VaR under the Base Case assumptions); and (b) that traders' asset allocation strategies should remain constant. Because we feasibly cannot have these traders go back and trade under different restrictions, we consider this to be the soundest approach to mimicking their intentions *ex post*.

Finally, we aggregate the adjusted positions of all the traders to reach a Trading Floor portfolio using the formula:

$$(10) \quad Q_i = \sum_{t=1}^{t=T} q_{i,t} ,$$

for all assets i , where $q_{i,t}$ is the size of the position of trader t in asset i , T is the number of traders participating in the Game, and Q_i is the Trading Floor's resulting aggregated position in asset i .

Once all asset positions are aggregated, we calculate the Trading Floor's annualized returns and annualized volatilities. Finally we determine the Trading Floor's Sharpe ratio using the formula:

$$(11) \quad S = R/\sigma ,$$

where S is the Sharpe ratio, R is the annualized return of the Trading Floor, and σ is the annualized volatility of that return.

This approach provides considerable benefits over other possible methods of adjusting positions after the fact. The first stems from a practical reality of risk management on large trading floors. From our discussions with risk managers, we gleaned that they are primarily concerned with an individual trader's overall risk exposure, not the asset-by-asset components of that total risk. By matching each trader's daily *portfolio* VaR, we are maintaining the integrity of the total risk figure.¹⁵ Another benefit is that because all assets are weighted by the same factor when bringing the portfolio risks in line under the new input assumptions, the asset allocation remains proportionally the same as the trader originally intended. So, regardless of what happens under the dynamic VaR constraints, *each trader's* asset allocation strategy remains intact.

¹⁵ Unavoidably, this means that the individual asset VaRs measured using dynamic inputs do not mirror their original counterparts; it is not possible to match both asset *and* portfolio VaRs using different input assumptions.

This does not imply, however, that the entire Trading Floor's asset allocation will be the same under each methodology. This is because those traders with positions in assets whose volatility estimates changed most (under the different methodologies) were compelled to adjust their portfolios most. A stylized example is to consider a two-trader Trading Floor, in which one trader has \$100 invested in equities and the other has \$100 invested in bonds under the Base Case assumptions, resulting in a Trading Floor asset allocation of 50%-50% between equities and bonds. If equity volatility falls by 50%, the equity trader will be able to increase his position to \$200 while maintaining a constant VaR. Assuming bond volatility and the equity-bond correlation remain constant, the Trading Floor's new asset allocation is now 67%-33% between equities and bonds, even though neither trader changed his individual asset allocation. Typically this is a desirable result, due to the strong correlation between falling market volatility (as with equities in this example) and rising asset prices (Rosenthal, 2006).

VaR Input Estimates

In order to calculate the single portfolio VaR statistic, estimates of current volatilities and correlations must be made. These correlation and volatility estimates feed into the VaR calculation and serve to estimate new VaR levels to measure risk. There are different approaches to choosing volatility and correlation inputs; we compare static, historically observed assumptions versus dynamic, rolling ones.

Original VaR Estimates

The Base Case assumptions used in the Game in 2006, 2007, and 2008 reflect long-term, historically observed correlations and volatilities. Individuals participating in the Game were constrained by VaR measurements evaluated using these metrics.

Dynamic Volatility Estimates

Both of the dynamic input schemes incorporate rolling, weighted volatility estimates. To achieve this, we first calculate squared daily log returns of each asset, beginning 360 days prior to the beginning of each trading period:

$$(12) \quad u^2 = \ln (P_1/P_0)^2 ,$$

where P_1 represents an asset's price on a given day, and P_0 represents that asset's price on the previous day. Each squared log return (u^2) represents the daily variance of that asset.

Next, our weighting scheme calculates a series of weights to apply to each day's asset variances. Weights decay exponentially in accordance with the following formula:

$$(13) \quad c_n = (1 - \lambda)(\lambda^n) ,$$

where c_n is the weight, n indicates how many days in the past a specific return occurred, and λ represents the exponential decay factor.¹⁶ The value of c_0 will always be the weight applied to the current day.

With the weights decaying exponentially among dates, we can calculate weighted daily variances by multiplying each day's weight by that day's variance for each asset. This serves to bias the resulting volatility toward recent daily volatilities. We then calculate a weighted annualized volatility (a_i) for each asset by summing the most recent 360 daily weighted variances, multiplying this by 252,¹⁷ and taking the square root of this result:

$$(14) \quad a_i = \sqrt{252 * \sum_{n=0}^{360} c_n u_{n,i}^2} ,$$

¹⁶ The specific decay factors (λ) are provided in Table 4.

¹⁷ The number 252 reflects the number of trading days in a year.

for all i . Each day, we arrive at a new weighted annualized volatility for each asset, as this estimate is rolling and incorporates new data on a daily basis.

Dynamic Correlation Estimates

To model the dynamic correlation estimates, we first calculate the log returns of each asset. Next, each day, we calculate 90-day rolling correlations of the returns of each of the 45 Trading Game asset pairs¹⁸ for the past 360 days. Then, for each asset pair, we exponentially weight these 360 90-day correlation estimates using the same lambda decay function¹⁹ as we used for volatilities. Finally, with these figures, we calculate a 360-day rolling weighted correlation ($b_{i,j}$) by summing the last 360 90-day weighted correlation estimates:

$$(15) \quad b_{i,j} = \sum_{n=0}^{360} c_n \rho_{n,i,j} ,$$

for all i and j . Here, i and j represent different assets; n represents the number of days before the current day; c_n represents the weighting factor for the 90-day correlation estimate ending on day n ; and $\rho_{n,i,j}$ is the ninety-day correlation of the returns of assets i and j spanning the period from n days in the past to $n + 90$ days in the past.

Overview of Methodologies 1 – 3

We run the Horse Race using three different input methodologies of measure the VaR constraints on the traders' positions. Each methodology applies a different set of input assumptions to the measurement of the traders' VaRs by using the volatility and correlation estimates given above. The decay variable (λ) in the exponential weighting scheme (for Methodologies 2 & 3) is the factor that defines how quickly the weights decay across a series of

¹⁸ 55 pairs in 2008 with the inclusion of VIX.

¹⁹ See Formula 13.

data. A λ closer to 0 causes the weights to decay more quickly, while a λ closer to 1 results in weights decaying more slowly. Each of the three methodologies is outlined below and summarized in Table 4.

Methodology 1

Methodology 1 replicates the Base Case assumptions imposed during the original Trading Games.

Methodology 2

Methodology 2 uses rolling, weighted volatilities and rolling, unweighted correlations. It weights volatilities with $\lambda = 0.944$ and correlations with $\lambda = 0.000$, putting all weight on the most recent 90-day correlation estimate.²⁰ RiskMetrics provides the precedent for weighting volatilities with a $\lambda = 0.944$ (Andersen *et al*, 2003).

Methodology 3

Methodology 3 uses rolling, weighted volatilities ($\lambda = 0.944$) and rolling, weighted correlations ($\lambda = 0.999$).²¹ For both Methodologies 2 & 3, a λ close to 1 distributes weight exponentially among past volatilities and correlations (“weighted” constraint). Due to the nature of the exponential decay function, it isn’t possible to have a λ equal exactly 1.

²⁰ This is represented by the following formula: $c_0 = 1$; $c_n = 0 \forall n \neq 0$, where c is the weight, and n indicates how many days in the past an observation occurred.

²¹ We decided against using any methodology with rolling, unweighted volatility constraints. Because VaR would be calculated with all weight on the previous day’s volatility, traders would have been able to take on enormous positions any day immediately following a day of very little movement in asset prices. This would have resulted in extreme adjusted positions, and we therefore deemed it an uninteresting constraint scheme to consider.

Table 4: Summary of Methodologies 1 – 3

Methodology	Volatility Estimate	Volatility λ	Correlation Estimate	Correlation λ
1	Static	-	Static	-
2	Rolling, Weighted	0.944	Rolling, Unweighted	0.000
3	Rolling, Weighted	0.944	Rolling, Weighted	0.999

IV. Results

Our results support our hypothesis that rolling, dynamic volatility and correlation assumptions, as used in Methodologies 2 & 3, are preferable in all markets to the static, historical inputs used in Methodology 1. In 2008, Methodologies 2 & 3 both resulted in final annualized returns that were significantly less negative than those of Methodology 1 *and* significantly less volatile.²² The same result was observed in 2007; Methodologies 2 & 3 performed better than Methodology 1, as they resulted in both higher annualized returns and lower volatilities. In 2006, Methodology 1 resulted in a negative annualized return, while Methodologies 2 & 3 attained positive annualized returns, but with slightly higher volatilities than Methodology 1. Dynamic Methodologies 2 & 3 tend to impose tighter risk limits compared to Methodology 1 when assets have negative returns and looser ones when assets' values increase.

2008 Results Discussion

September 8 to December 5 of 2008 was a period of great uncertainty in the financial markets. Over this time horizon the Dow Jones Industrial Average and the S&P 500 experienced annualized returns of -64% and -74%, respectively. At the same time, debt instruments varied widely in their returns. On an annualized basis, 10-year US Treasury prices rose 36%, while Baa Corporate Bond prices fell 38%. One of the few consistencies during the autumn of 2008 is the extreme volatility across all asset classes. In all assets aside from the Japanese Yen, our rolling volatility scheme limited traders' positions more than the static volatilities used in the Base Case.²³ As expected, overall Trading Floor positions were reduced under Methodologies 2 & 3,

²²Throughout this section, all references of returns and the volatility of those returns are at the Trading Floor level.

²³ See Table 6.

limiting the Floor's exposure to the market's considerable volatility and resulting in final P&Ls much less negative than that of the Trading Floor under Methodology 1.

Under Methodology 1, the annualized return of the Trading Floor in 2008, comprised of seventy-nine individual traders, was -30.238%; additionally, the annualized volatility of these returns was 2,000%. In total, the floor recorded a loss of \$76,213,133 and a Sharpe ratio of -15.12, as shown in Table 5. The Floor's largest losses were seen in the equity market, as it realized a loss of \$63,096,506 from its equity positions²⁴. Though the Trading Floor achieved a profit of \$3,257,956 from its positions in debt, its large exposure to equities dominated the final P&L figure and was primarily responsible for its 2008 losses. Clearly, the static volatilities used to measure VaR under Methodology 1 were inadequate, given the observed market volatility in 2008. This methodology based its portfolio limits on outdated volatility and correlation assumptions, and it was thus unfit to constrain traders in this market environment.

The Trading Floor fared better in 2008 under the dynamic assumptions of Methodology 2 as compared to the Base Case performance. The rolling volatilities were largely more restrictive across assets, leading to a final Trading Floor loss of \$28,224,328, quite less than that under Methodology 1. Furthermore, the Trading Floor recorded an annualized return of -347%, an annualized volatility of 54%, and a Sharpe ratio of -6.47 under Methodology 2. The Trading Floor loss attributable to the equity markets was \$28,151,513, notably less than the equity losses sustained by the Trading Floor in Methodology 1. Also, the Floor remained profitable in the debt markets, gaining \$2,433,314. As evidenced in Table 6, the rolling volatility estimates in this methodology were better able to reflect the uncertain environment of the 2008 marketplace.

²⁴ Tables 11 – 13 provide a detailed breakdown of traders' P&L by asset for each year under all three methodologies.

Similarly, the results of Methodology 3 represent a large improvement over those of Methodology 1. The Trading Floor realized a final loss of \$28,475,020, with an annualized return of -354%, an annualized volatility of 56%, and a Sharpe ratio of -6.32. Again, these figures represent an improved P&L with lower annualized volatility than that of the Trading Floor under the Base Case assumptions. As in Methodology 2, the Trading Floor was more limited in its exposure to the equity markets, experiencing a final, equity-specific loss of \$28,758,264. Additionally, the Trading Floor realized a profit in the debt markets amounting to \$2,704,558.

As previously stated, Methodologies 2 & 3 estimated rolling volatilities that were considerably higher than the Base Case volatilities of Methodology 1. For example, the Trading Floors of Methodologies 2 & 3 were constrained by a VaR measured using an annualized volatility of 62%²⁵ for the S&P 500, while Methodology 1 used an annualized volatility estimate of only 13%. Similar results are seen for the DJIA, as our rolling assumptions estimated an annualized volatility of 55%. The static inputs again limited traders with an annualized volatility of only 13%. The rolling volatilities proved to be more accurate input assumptions, as the S&P 500, over the Trading Game period, returned an annualized loss of 74% and the DJIA witnessed a similar loss of 64%. Moreover, the rolling volatility estimates for commodities were considerably higher than their Base Case counterparts of Methodology 1. While oil and gold were constrained using respective volatility estimates of 32% and 19% in Methodology 1, positions in these same assets were limited using volatility estimates of 81% and 34% respectively in Methodologies 2 & 3. This again led to a more favorable Trading Floor outcome, as oil and gold actually realized respective annualized returns of -98% and -22% over this period.

²⁵ As in Tables 6, 8, and 10, quoted volatilities for Methodologies 2 & 3 in this section reflect the volatility calculation for that asset on the final day of trading for that year's Game, as that date encompasses all data over that time period.

Table 5: 2008 Horse Race Results

Methodology	1	2	3
Annualized Return	-30,238%	-347%	-354%
Annualized Volatility	2,000%	54%	56%
Sharpe	-15.12	-6.47	-6.32
P&L	-\$76,213,133	-\$28,224,328	-\$28,475,020

Table 6: 2008 Static and Rolling Volatility Estimates, Observed RoR, and Portfolio Weights from 9/8/2008 to 12/5/2008

Rolling Volatilities and Portfolio Weights are as of the last day of the Trading Game period²⁶

		Annualized				Portfolio Weight		
		Static	Rolling	Difference	RoR	1	2	3
<i>Less constrained by rolling estimates</i>	\$/Yen	22%	19%	-3%	78%	7%	13%	12%
	2yr UST	2%	2%	0%	11%	2%	4%	3%
	Baa Corp Bond	5%	10%	5%	-38%	16%	19%	19%
	10yr UST	7%	14%	7%	36%	11%	10%	9%
	\$/Euro	9%	16%	7%	-36%	2%	-4%	-1%
	Gold	19%	34%	16%	-22%	28%	33%	31%
	Nasdaq	36%	62%	26%	-79%	7%	5%	6%
	VIX	96%	138%	42%	3848%	2%	1%	2%
	DJIA	13%	55%	43%	-64%	19%	19%	17%
	Oil	32%	81%	48%	-98%	6%	6%	6%
<i>More constrained by rolling estimates</i>	S&P500	13%	62%	49%	-74%	-1%	-6%	-3%
						Equities	25%	18%
						Debt	30%	33%
						Commodities	34%	39%
						Currency	10%	9%

²⁶ We chose to show rolling volatilities as of the last day of the Trading Game period so as to incorporate all the data from the actual period and likewise to show portfolio weights as of the period's last day in order to depict a final "snapshot" of the asset allocation of the Trading Floor.

2007 Results Discussion

The Trading Game period in 2007, from September 10 to November 29, was a prototypical mild rise-and-fall market. The annualized return for the period of the broad-based S&P 500 equity index was a mere 3%, in line with the rate of inflation for 2007. Other asset classes provided more lucrative returns during this period, however. Oil and gold rose at annualized rates of 84% and 51%, respectively. In this tepid market, commodities notwithstanding, rolling estimates did not differ much from the Base Case assumptions. Despite this, the Trading Floor constrained by a VaR measured with either of the alternate dynamic methodologies still fared better than it did under the Base Case, achieving both higher rates of return and lower volatilities.

Under the Base Case estimates, the Trading Floor, encompassing fifty-three individual traders, gained a cumulative \$778,122 during this time period. As seen in Table 7, the annualized return for the Trading Floor was 11%, with an annualized volatility of 29%, resulting in a Sharpe ratio of 0.38. This profit resulted primarily from large gains in gold (\$2,664,488) and the Japanese Yen (\$1,474,970), though the Trading Floor lost \$1,083,247 on its debt position, stemming from its short positions in a rising credit market.

The Trading Floor performed best when constrained by Methodology 2, which was characterized by rolling, weighted volatilities and rolling, unweighted correlations. Using this scheme, the Trading Floor achieved \$1,658,005 in profits, with an annualized return of 18%, an annualized volatility of 25%, and a Sharpe ratio of 0.72. Again, this was the result of successful bullish bets on gold (\$2,748,787) and the Yen (\$1,014,434). Also, under this input methodology, the Trading Floor trimmed its bearish bond positions somewhat and only lost \$830,707 on its debt exposure.

The Trading Floor under Methodology 3 witnessed similar return and volatility levels as under Methodology 2. Using rolling weighted volatilities and correlations, the Trading Floor profited \$1,544,009, representing an annualized return of 17%, an annualized volatility of 26%, and a resulting Sharpe ratio of 0.65. As with the other two schemes, this stemmed primarily from gains on gold (\$2,815,793) and the Yen (\$1,002,118). Once again, though, it sustained losses from its overall short debt position, losing \$938,960 on that asset class.

In this rather neutral 2007 market, the dynamic schemes were more successful than the Base Case, primarily due to a *mid-period* correction. The rolling estimates for the equity indices were quite low (and thus less restrictive) during the earlier part of this trading period, when the equity market was performing well.²⁷ This is because they were reflecting the positive market conditions of late 2006 and early 2007. Soon, however, this rapid early autumn rise began to register in the rolling equity volatilities, causing them to increase, which in turn forced the traders to come closer to home and reduce their positions just as the equity market began its late autumn descent.²⁸

²⁷ The S&P 500 rose at an impressive 57% annualized rate from September 10 to October 31 of 2007 before falling precipitously throughout the rest of the Trading Game period.

²⁸ Table 8 illustrates these elevated, late-period equity volatilities.

Table 7: 2007 Horse Race Results

Methodology	1	2	3
Annualized Return	11%	18%	17%
Annualized Volatility	29%	25%	26%
Sharpe	0.38	0.72	0.65
P&L	\$778,122	\$1,658,005	\$1,544,009

Table 8: 2007 Static and Rolling Volatility Estimates, Observed RoR, and Portfolio Weights from 9/10/2007 to 11/29/2007

Rolling Volatilities and Portfolio Weights are as of the last day of the Trading Game period²⁹

		Static	Rolling	Difference	Annualized RoR	Portfolio Weight		
						1	2	3
<i>Less constrained by rolling estimates</i>	Oil	38%	24%	-13%	84%	19%	11%	9%
	\$/Euro	10%	6%	-4%	26%	-3%	-3%	-1%
	2yr UST	2%	3%	1%	8%	-308%	-137%	-119%
	Baa Corp Bond	5%	6%	1%	1%	107%	55%	52%
	10yr UST	7%	8%	1%	14%	-99%	-46%	-43%
	DJIA	17%	19%	1%	2%	108%	64%	58%
	S&P500	18%	21%	3%	3%	68%	35%	34%
	Nasdaq	30%	33%	3%	11%	68%	39%	37%
<i>More constrained by rolling estimates</i>	Gold	17%	22%	6%	51%	161%	87%	77%
	\$/Yen	10%	18%	8%	11%	-19%	-5%	-4%
						Equities	244%	138%
						Debt	-300%	-128%
						Commodities	179%	99%
						Currency	-22%	-8%
								-5%

²⁹ We chose to show rolling volatilities as of the last day of the Trading Game period so as to incorporate all the data from the actual period and likewise to show portfolio weights as of the period's last day in order to depict a final "snapshot" of the asset allocation of the Trading Floor.

2006 Results Discussion

September 6 to November 29, 2006 was a largely favorable market for investors. The S&P 500 rose at an annualized rate of 30% and the benchmark 10-year US Treasury bond advanced at an annualized rate of 11%. Daily asset price volatilities and correlations between asset pairs were mild, as asset values, with the exception of commodities, marched higher throughout the period. It is in this type of market that we expect VaR using dynamic inputs to limit traders' positions less and allow them to take on more risk.

Under the Base Case, the Trading Floor, encompassing fifty-four individual traders, lost a cumulative \$693,879 during this time period, as illustrated in Table 9. The annualized return for the Trading Floor was -5%, with an annualized volatility of 12%. This resulted in a Sharpe ratio of -0.39. Such a result is perhaps surprising considering that it occurred during generally positive macroeconomic conditions. Breaking down this overall P&L by asset class shows that while the Trading Floor was able to profit \$2,226,643 from the rising equity market, it more than offset these gains with its losses in gold (-\$2,594,400) and bonds (-\$819,462).

When constrained by Methodology 2, however, the Trading Floor performed better, making an overall profit of \$249,231. This represents an annualized return of 3%, an annualized volatility of 17%, and a Sharpe ratio of 0.20. The looser VaR constraints calculated using this input methodology resulted in larger losses from the Trading Floor's gold position (-\$2,901,553), and bond position (-\$1,571,626). The lower volatility estimates for equities, though, allowed the Trading Floor to take large stakes in this asset class, and its vast equities profit of \$4,261,125 resulted in an overall profit for the Trading Floor during this period. Table 10 illustrates that the dynamic volatility assumptions were lower than those of the Base Case for *all* assets, allowing traders to increase all of their positions in this bull market.

A similar result is observed using Methodology 3. Under these input assumptions, the Trading Floor profited \$539,905 during the period, representing an annualized return of 6%, an annualized volatility of 16%, and a resulting Sharpe ratio of 0.34. As with Methodologies 1 & 2, the Trading Floor lost large amounts on gold (-\$2,807,611) and bonds (-\$1,507,016). Again, though, this methodology with rolling, weighted volatilities and correlations allowed the Trading Floor to exploit a tranquil equity market and make large equity positions that netted the Trading Floor \$4,258,267.

These considerably different results between the Base Case and the dynamic input methodologies stem from the manner in which the individual traders were constrained. As in the 2007 and 2008 periods, traders tended to dedicate a large proportion of their risk allowance to the equity asset class. Consequently, it is particularly necessary for equity volatility estimates to closely reflect the actual current market conditions. In 2006, as is typical of favorable equity markets, volatilities for all the equity indices were quite low. On the last day of the 2006 Trading Game period, the 360-day rolling volatilities for the Nasdaq, S&P 500, and DJIA were only 11%, 8%, and 7% respectively. These figures are substantially lower than the respective Base Case volatility assumptions of 31%, 18%, and 18%, resulting in looser VaR constraints under dynamic Methodologies 2 & 3. As expected, these looser VaR constraints coincided with a rising and frothy equity market. During this time period, the Nasdaq, S&P 500, and DJIA rose at annualized rates of 50%, 30%, and 30% respectively.

While the rolling volatilities of all assets were lower at this time than the static volatility estimates, equities' volatilities in particular declined most dramatically. This meant that, while the Trading Floor increased all of its positions, it was also incentivized to dedicate a larger proportion of its assets to equities because each unit of equity traded now "consumed" relatively

less of each trader's risk allowance. Table 10 provides a snapshot of the Trading Floor's final portfolio allocation showing this result. Under the higher Base Case equity volatility estimates, the Trading Floor allocated 62% of its portfolio to equities. However, using dynamic input assumptions, the Floor allocated 150% and 156% of its portfolio to the equity asset class under Methodologies 2 & 3, respectively.³⁰ The direct implication of these figures is that the Trading Floor took a more bullish position in equities precisely when this market was performing best. That is to say, without individual traders necessarily having a distinct view on future equity performance, the more robust volatility input scheme compelled the Floor to allocate proportionally more portfolio weight to the equity market, where volatilities were most historically low.³¹

Table 9: 2006 Horse Race Results

Methodology	1	2	3
Annualized Return	-5%	3%	6%
Annualized Volatility	12%	17%	16%
Sharpe	-0.39	0.20	0.34
P&L	-\$693,879	\$249,231	\$539,905

³⁰ These percentages exceed 100% because short positions are allowed.

³¹ As noted before, the dynamic input schemes do not compel individual traders to invest proportionally more in the equity market, but rather allow those traders with equity stakes (in this particular situation) to take on larger positions, thus causing the overall Trading Floor's equity proportion to increase.

Table 10: 2006 Static and Rolling Volatility Estimates, Observed RoR, and Portfolio Weights from 9/6/2006 to 11/29/2006

Rolling Volatilities and Portfolio Weights are as of the last day of the Trading Game period³²

		Annualized				Portfolio Weight		
		Static	Rolling	Difference	RoR	1	2	3
<i>Less constrained by rolling estimates</i>	Nasdaq	31%	11%	-20%	50%	25%	57%	59%
	S&P500	18%	8%	-11%	30%	30%	72%	72%
	DJIA	18%	7%	-11%	30%	7%	22%	25%
	Oil	38%	28%	-10%	-22%	19%	30%	37%
	\$/Yen	10%	5%	-5%	1%	4%	15%	10%
	\$/Euro	10%	5%	-5%	10%	5%	11%	8%
	Baa Corp Bond	7%	3%	-4%	12%	-2%	-16%	-18%
	Gold	16%	13%	-3%	-7%	20%	45%	54%
<i>More constrained by rolling estimates</i>	2yr UST	2%	1%	-1%	1%	-168%	-262%	-277%
	10yr UST	4%	4%	-1%	11%	-39%	-74%	-69%
						Equities	62%	150%
						Debt	-209%	-352%
						Commodities	38%	75%
						Currency	9%	26%
								18%

Comparison of Results of Methodology 2 vs. Methodology 3

With its unweighted correlations more quickly reflecting changing market environments, Methodology 2 outperformed Methodology 3 in both the hostile market of 2008 and the rather neutral 2007 market. In both periods, Methodology 2 resulted in a higher return and lower volatility than Methodology 3. Methodology 2 placed all weight on the most recent set of 90-day correlation estimates and thus was better suited to capture the immediate effects of changes in correlations. This result is in line with our hypothesis that the most effective input scheme is the one that most accurately reflects the current state of the markets.

However, in both 2008 and 2007, the difference between the results of Methodology 2 and Methodology 3 was quite small. During the 2008 Game period, the changes in volatilities

³² We chose to show rolling volatilities as of the last day of the Trading Game period so as to incorporate all the data from the actual period and likewise to show portfolio weights as of the period's last day in order to depict a final "snapshot" of the asset allocation of the Trading Floor.

were so great that they overwhelmed the changes in correlations.³³ This considerable increase in volatility in all asset classes led to sharply reduced positions under either Methodology 2 or 3.

As a result, the effect of the different correlation weighting schemes was relatively inconsequential. During the 2007 Game period, correlations (and volatilities) were not very dynamic.³⁴ Consequently, the difference between the two correlation weighting schemes did not have a great impact on the computed correlation estimates, and thus the VaR constraints under both methodologies were quite similar.

The results from the 2006 Trading Game period, however, are markedly different than those discussed above. In this market, Methodology 3, which is *slower* to incorporate new correlation information, outperformed the more responsive Methodology 2. This outcome stems from an intrinsic reality of using VaR in a portfolio that allows short positions. The negative relationship between correlations and returns is robust and was observed in this particular market. During the second half of the 2006 Trading Game period, the benchmark 10yr US Treasury rose at an annualized rate of 22%, while its correlation with the S&P500 fell from 25% to -1%.³⁵ The dynamic input schemes captured this falling correlation, and the VaR constraints on these assets consequently loosened, allowing traders to take bigger positions in them. This is how we expected the dynamic inputs to behave, and if the Trading Floor were *long* in these assets, it indeed would benefit from being able to take on larger stakes just as these assets' values are rising.

³³ To check this “dominance” of the volatility inputs, we adjusted traders’ positions using dynamic volatilities and original *static* correlation estimates. As expected, these results were very similar to the results under Methodologies 2 & 3 in 2008, confirming our assumption that 2008 volatility inputs largely overshadowed the effects of correlation estimates, either static or dynamic.

³⁴ See Chart 11.

³⁵ These figures were calculated using 90-day correlation estimates from 10/16/2009 through 11/29/2006.

In 2006, however, the Trading Floor had very large *short* positions in the debt market, and these lower correlation estimates allowed its traders to go even shorter amid a *rising* debt market. That is an inherent issue with dynamic VaR inputs: if the trader is short a rising market, volatilities and correlations will tend to fall and the trader will be allowed to increase the magnitude of his short position and in essence become “more wrong.”³⁶ This is why Methodology 3 outperformed Methodology 2 in this favorable market; by design, Methodology 3’s weighting scheme was slower to capture these falling correlations, and thus it was not as quick to allow the traders to increase the size of their misguided short debt positions.³⁷ These results all illustrate the importance of context in risk management. As the recent financial crisis has shown, it is necessary to understand the current pulse of the market and not just apply a single model. It is in this sense that risk management is an art as well as a science.

³⁶ It should be noted that this is not an issue if a trader is long a falling market, because in that situation dynamic correlation and volatility estimates will tend to increase, thus compelling the trader to cut his positions and minimize his losses.

³⁷ Further support of this arose from our “check” (see footnote 31) using dynamic volatility and original *static* correlation assumptions. In essence, static correlation estimates are the *slowest* to incorporate new data (as they do not incorporate new data at all). Thus, in 2006, this set of dynamic volatility and static correlation assumptions outperformed Methodology3 for the same reason Methodology 3 outperformed Methodology 2.

Table 11: 2008 Trading Floor P&L by Asset
(\$000s)

	2yr UST	10yr UST	Baa Corp Bond	S&P500	Nasdaq	DJIA	Gold	Oil	\$/Euro	\$/Yen	VIX
Methodology 1	\$994	\$917	\$1,347	-\$22,635	-\$13,151	-\$27,311	-\$4,043	-\$14,599	-\$2,696	\$1,909	\$7,711
Methodology 2	\$825	\$785	\$823	-\$9,544	-\$5,782	-\$12,826	\$34	-\$5,834	-\$1,527	\$1,392	\$5,810
Methodology 3	\$968	\$949	\$787	-\$10,083	-\$6,008	-\$12,667	\$390	-\$6,110	-\$1,502	\$1,326	\$5,877

Table 12: 2007 Trading Floor P&L by Asset
(\$000s)

	2yr UST	10yr UST	Baa Corp Bond	S&P500	Nasdaq	DJIA	Gold	Oil	\$/Euro	\$/Yen
Methodology 1	-\$455	-\$524	-\$104	\$183	-\$2	\$504	\$2,664	\$46	\$114	\$1,475
Methodology 2	-\$620	-\$232	\$22	\$428	-\$475	\$772	\$2,749	-\$169	\$512	\$1,014
Methodology 3	-\$687	-\$301	\$49	\$334	-\$412	\$739	\$2,816	-\$100	\$485	\$1,002

Table 13: 2006 Trading Floor P&L by Asset
(\$000s)

	2yr UST	10yr UST	Baa Corp Bond	S&P500	Nasdaq	DJIA	Gold	Oil	\$/Euro	\$/Yen
Methodology 1	-\$407	-\$345	-\$68	\$818	\$1,340	\$69	-\$2,594	\$340	\$257	-\$44
Methodology 2	-\$669	-\$676	-\$227	\$1,654	\$2,457	\$150	-\$2,902	\$664	\$382	\$29
Methodology 3	-\$694	-\$628	-\$186	\$1,664	\$2,444	\$150	-\$2,808	\$731	\$361	\$39

V. **Conclusion**

Our results provide sufficient evidence in support of our hypothesis of the preferability, in all markets, of incorporating more recent, rolling data to formulate the statistical input assumptions of the VaR framework. While the extant research currently abounds with work focusing on replacing or fundamentally altering the VaR model, less attention is granted to improving the quality of the model's basic underpinnings, which drive VaR in its current form. The consistency and robustness of our results support our belief that research striving to improve the quality of the statistical inputs of the VaR model is a practical means of advancing a risk management tool that seems unlikely to fade away soon. In particular, the future of risk management at financial firms looks certain to be increasingly driven by external factors due to impending regulatory overhaul. The current public distrust of Wall Street complexity and opacity seems to bode well for the continued preference for the "seductively simple" VaR framework, if only as a high-level risk barometer.

One particular area our research leaves open for future study deals with VaR inputs' inherent inequitable handling of short and long positions. While the behavior of correlations and volatilities aids long positions in both bull and bear markets, they hinder the performance of short positions in both market environments. In a bear market, a trader's short position will be profitable; however, dynamic VaR estimates will inevitably force the trader to decrease his positions (in an absolute sense) over time, thereby reducing his profitability as the markets continue to fall.³⁸ Similarly, in a bull market, the dynamic estimates, now likely less constrictive, will compel traders with short positions to increase these positions and incur enlarged losses as the markets rise. This is only one example of the interesting interplay that

³⁸ Please note that the dynamic statistical inputs were more effective in the bear market of 2008, because the traders were *long*, rather than short, in this environment (see Table 6).

exists between the VaR model and the statistical input assumptions upon which it is dependent; an increased appreciation for the benefits that can be achieved by improving these inputs will hopefully prompt further research in this area.

References

- Andersen, T. G., Bollerslev, T., Christoffersen, P.F., & Diebold F.X. (2005). Practical Volatility and Correlation Modeling for Financial Market Risk Management. In Carey, M., & Stulz, R. (Ed.), *Risks of Financial Institutions* (pp. 1-40). Chicago: University of Chicago Press for National Bureau of Economic Research.
- Andersen, T. G., Bollerslev T., Diebold F. X., & Labys, P. (2003). Modeling and Forecasting Realized Volatility. *Econometrica*, 71(2), 579-625.
- Bangia, A., Diebold, F. X., Schuermann, T., and Stroughair, J. (2002). Modeling Liquidity Risk, with Implications for Traditional Market Risk Measurement and Management. In Figlewski, S., & Levich, R. (Ed.), *Risk Management: The State of the Art* (pp. 1-13). Amsterdam: Kluwer Academic Publishers.
- Beder, T. S. (1995). VAR: Seductive but Dangerous. *Financial Analysts Journal*, 51(5), 12-24.
- Bogle, J. C. (2008). Black Monday and Black Swans. *Financial Analysts Journal*, 64(2), 30-40.
- Berkowitz, J., & O'Brian, J. (2002). How Accurate are Value-At-Risk Models at Commercial Banks? *Journal of Finance*, 57(3), 1093-1111.
- Gilli, M., & Kellezi, E. (2006). An Application of Extreme Value Theory for Measuring Financial Risk. *Computational Economics*, 27(2), 207-228.
- Jorian, P. (1996). Risk²: Measuring the Risk in Value at Risk. *Financial Analysts Journal*, 52(6), 47-56.
- Krokhmal, P., Jonas P., & Uryasev, S. (2001). Portfolio Optimization with Conditional Value-at-Risk Objective and Constraints. *Journal of Risk*, 4(2).
- Lucas, A. (2000). A Note on Optimal Estimation from a Risk Management Perspective under Possibly Misspecified Tail Behavior. *Journal of Business and Economic Statistics*, 18(1), 31-39.
- Nocera, J. (2009). Risk Mismanagement. *The New York Times*. Retrieved July 27, 2009, from http://www.nytimes.com/2009/01/04/magazine/04riskt.html?_r=1&scp=1&sq=risk%20mismanagement&st=cse.
- Patton, A. (2006). Copula-based Models for Financial Time Series. In Andersen, T. G., Davis, R. A., Kreiß, J. P., & Mikosch, T. (Ed.), *Handbook of Financial Time Series* (pp. 767-85). Berlin: Springer-Verlag.
- Pine Cook Capital, LLC. (2008). A Study of Fat-Tail Risk. Retrieved November 11, 2009, from <http://www.cookpinecapital.com/pdf/Study%20of%20Fat-tail%20Risk.pdf>.

- Romano, C. (2002). Applying Copula Function to Risk Management. *La Sapienza*. University of Rome, working paper. Retrieved December 1, 2009, from www.gravitascapital.com/Research/Risk/Applying%20Copula%20Function%20to%20Risk%20Management.pdf.
- Richardson, M., & Smith, T. (1993). A Test for Multivariate Normality in Stock Returns. *Journal of Business*, 66(2), 295-321.
- Rosenthal, S. J. (2006). Volatility Rising: The Growth of an Alternative Asset Class. *Credit Suisse*. Retrieved November 10, 2009, from <http://www.euromoney-yearbooks.com/images/143//downloads/creditsuissefinal.pdf>.
- Salvadori, G., & De Michele, C. (2007). On the Use of Copulas in Hydrology: Theory and Practice. *Journal of Hydrologic Engineering*, 12(4), 369–380.
- Sheikh, A. Z., & Qiao, H. (2009). Non-normality of Market Returns: A Framework for Asset Allocation Decision-making. *J.P. Morgan Asset Management*. Retrieved October 20, 2009, from www.jpmorgan.com/directdoc/nonnormality.pdf.
- Taleb, N. (2007). *The Black Swan: The Impact of the Highly Improbable*. New York: The Random House Publishing Group.
- Yi, W., & Bier, V. M. (1998). An Application of Copulas to Accident Precursor Analysis. *Management Science*, 44(12), 257-270.