Approaches for Avoiding Margin Call Situations Through Risk-Management Automation

Chad Dau, Matthew Groch

Practicum in Analytical Finance
Spring 2005

Abstract

BlueSky Financial* is an online retail trading brokerage that must manage the risk associated with individual trading accounts exceeding their respective buying power and entering margin call situations. This paper presents alternative approaches for managing this risk and explores some of the performance issues that must be considered as a part of any practical implementation within an automated risk-management back office.

* The actual name of the brokerage has been disguised for confidentiality purposes. We would like to thank Mr. Brent Stark for his sponsorship and we would also like to thank Professor Ravi Jagannathan and Professor Ernst Schaumburg for their help and guidance.

1. Introduction

BlueSky Financial is an online brokerage firm that provides stock and option trading services to retail customers. Retail customers manage cash, stock, and option positions and execute trades via individual trading accounts. Each account's trading activity is limited by its buying power, where buying power is measured as a multiple of the equity in an account. These multiples are determined by SEC and exchange rules and differ based on the type and characteristics of each customer account.

Like other online retail brokerages in the same market space, BlueSky Financial must constantly monitor these accounts and take action when an account's buying power reaches zero, or in other words, when an account enters a margin call situation. Margin call situations are risky for the firm because of the chance that the account will default and force the firm to assume any financial losses due to the equity in the account having been effectively wiped out. More than that, however, customer accounts that enter margin call situations multiple times are liable to be closed by the brokerage's clearing firm, thereby eliminating that revenue stream for the brokerage.

Margin call risk for brokerages must be managed from both a price movement as well as an open order perspective. With respect to price movements, there will likely exist certain customer accounts with asset portfolios that are heavily leveraged and particularly sensitive to movements in the overall market. Under certain market conditions, these types of accounts have the potential of having their buying power driven below zero in a short amount of time.

From an open order perspective, a customer account that maintain a relatively high number of open orders has a particularly higher chance of entering a margin call situation. For example, imagine an account in which a customer issues a large number of limit orders for various stocks and options, perhaps as a function of running several concurrent trading strategies. Given this example, a large enough move in the overall market might cause enough of these limit orders to execute with the end result, again, being an account with its buying power driven below zero in a short amount of time.

In an effort to mitigate the chances of margin call situations across its customer accounts, BlueSky Financial implemented specific rules in its back office trading system: their trading system automatically assumes that any limit order placed through the system will be filled and, as such, the buying power of customer accounts is debited accordingly upon submission of an order. Per limit order placed, buying power remains debited by an amount equal to the value of the order unless the order is explicitly canceled by the customer. While this solution certainly addresses the open order aspect of the firm's margin call risk, it does not address the price movement aspect of margin call risk and furthermore, it reduces the flexibility of the firm's customers and potentially has an impact on the overall revenues generated by the firm through customer transaction fees.

This paper briefly discusses the details of BlueSky Financial's current solution and then explores alternative approaches that address both the price movement aspect and the open order aspect of managing margin call risk. The paper describes how these alternative approaches are modeled and which of the approaches is a "best-fit" for the brokerage based on implementation-performance considerations. Building upon this discussion, the paper then outlines a conceptual framework for how this "best-fit" solution could be implemented into an automated back office trading/risk-management system.

2. Approaches To Managing Buying Power

Overview of Current Approach

When a new order[1] is submitted to the BlueSky Financial system, it remains an open order until either its limit price is hit due to movement in the market, the order expires (i.e. the order is specified as only being valid for the current trading day), or the order is cancelled by the customer. Technically, an account's buying power would not be immediately impacted by open orders since open orders are simply orders that have not yet executed. Rather, buying power would be debited (or credited, depending on whether a given order was on the buy or sell side of a transaction) only upon receipt of a fill (or partial fill) notification for a given open order.

[1] For the purposes of this paper, the term 'order' will imply limit orders, where limit orders are orders to buy or sell an asset at a certain price or better. Buy orders execute at or below a specific limit price and sell orders execute at or above a specified price. If the market price of an asset does not reach the specific limit price, the limit order will not execute. From a discussion standpoint, this paper will focus on limit orders since a) market orders execute immediately at the current market price, and b) stop orders are effectively modeled by limit orders as they become market orders at the time their stop price is reached by the market.

However, from a risk perspective, managing buying power via the technical standpoint is not a perfect solution. A customer could easily submit an arbitrary number of orders to the system whose combined value exceeds the buying power of an account and, upon execution, could potentially induce a margin call situation.

In order to mitigate this situation, BlueSky Financial has implemented rules whereby the automated trading system treats all orders entered into the system as if they are *guaranteed* to execute at their specified limit price. Restating this in a slightly more mathematical bent, the system assumes the probability of any order entered into the system being executed at its limit price as 100%. As each order is placed, account buying power is immediately updated to reflect a change in value equal to the product of the order size and limit price:

$$\Delta BP = size \times price$$

This solution effectively disallows a customer from placing orders that could potentially bring buying power down to zero (or less). However, the main drawback of this solution is that it severely limits the flexibility of an account in terms of a customer's ability to implement various trading strategies. The solution does not attempt to account for the fact that orders, as specified by customers, will actually have execution probabilities less than 100%. Furthermore, the solution also does not attempt to account for dependence within a set of two or more open orders, where dependence refers to statistical correlations in the rise and fall of asset prices. In other words, this solution does not provide buying power 'relief' in cases where there are sets of negatively correlated open orders associated with an account.

All told, this lack of flexibility and restriction on the placement of additional orders may result in customer frustration and lower retention of customer accounts within the brokerage. The highly conservative nature of this approach, by reducing the number of allowable trades, likely results in lower trade-based commissions than is potentially possible at an approximately equivalent level of risk for the brokerage.

Side	Ticker	Size	Price	Cost	BP
					5,000.00
Buy	AAPL	100	35.00	3,500.00	1,500.00

Figure 1

Figure 1 illustrates an example of this solution in action. Initially, a customer account has $5,000 in buying power. For the sake of this example, assume the current price of AAPL is $38 per share. After the customer submits an order to buy 100 shares of AAPL at a limit price of $35 per share, the system immediately debits the account buying power by 100 x $35 = $3,500, leaving $1,500 in buying power remaining. Next, imagine that the customer attempts to submit an order to buy 100 shares of SBUX at $50 per share for a total cost of 100 x $50 = $5,000. The system would reject this order since the account does not have a sufficient level of buying power. Furthermore, also assume the current price of SBUX is $90 per share, thus making the probability of that particular order being executed highly unlikely in any reasonable amount of time. Even so, the system will still reject the order.

Finally, assume the customer would like to construct a partial hedge against the AAPL order with a purchase of 100 shares of GLD at $35 per share with the current price at $40 per share. Even though there is a negative correlation between the two orders, the system does not take this into account and rejects this order as well.

Alternative Approach Disregarding Dependence in Price Movements

The first alternative approach this paper explores accounts for the fact that the probability of an order being executed is not necessarily 100%. However, this alternative still does not account for dependence within a set of two or more open orders, or rather, this alternative assumes that probabilities of order execution are independent (i.e. price movements of assets specified on open orders are independent of one another). With this approach, the expression for change in buying power becomes:

$$\Delta BP = size \times price \times \Pr(execution)$$

The interesting component of this expression is the probability of the order being executed. For the order to execute, the only event that must occur is the price of the asset specified by the order must hit or exceed the limit price specified by the order. Under an assumption of a Gaussian (i.e. normal) distribution of asset returns, the probability of the order being executed is described by a cash-or-nothing option[2], where the strike price of the option is the limit price specified by the order. Orders on the buy side are described by cash-or-nothing puts under the assumption that current market prices are always higher than limit prices for the buy side of a transaction. Orders on the sell side are described by cash-or-nothing calls under the assumption that current market prices are always lower than limit prices for the sell side of a transaction:

$$\Delta BP = size \times price \times e^{-r(T-t)} N(-d_2),\ Buy\ Side$$

$$\Delta BP = size \times price \times e^{-r(T-t)} N(d_2),\ Sell\ Side$$

$$d_2 = d_1 - \sigma\sqrt{T-t}$$

$$d_1 = \frac{\ln(\frac{S_t}{K}) + (r - \delta + \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}$$

The time to expiration for these cash-or-nothing options is a parameter that would be specified by the brokerage, where the brokerage would determine in advance to what time horizon (i.e. how far out into the future) margin call risk should be managed.

Side	Ticker	Size	Price	σ	(T-t)	Pr	ΔBP	BP
								5,000.00
Buy	AAPL	100.00	35.00	49.17%	0.00274	99.6459%	3,487.61	1,512.39
Buy	SBUX	100.00	50.00	29.63%	0.00274	8.6097%	430.48	1,081.91

Figure 2

[2] Given the assumption that prices of assets specified by orders follow geometric Brownian processes as described by $\frac{dS(t)}{S(t)} = \alpha\, dt + \sigma\, dZ(t)$, where *S(t)* is the current price of an asset, α is the expected return of the asset, σ is the volatility of the asset, and *Z(t)* is a normally-distributed random variable that follows a Brownian motion process.

Figure 2 illustrates an example of this alternative approach in action. As in the previous example, a customer account has $5,000 in buying power initially. In this example, the current price of AAPL is $38 per share and the current price of SBUX is $55 per share. Both AAPL and SBUX have a zero dividend policy ($\delta = 0$) and the risk-free rate is 3%. Additionally, assume the brokerage has specified 24 hours as the risk management time horizon (1/365 = 0.00274).

First, the order for 100 shares of AAPL is entered into the system. With a volatility of 49.17%, the probability of the order executing is approximately 99.65% thereby debiting buying power by $3,487.61 (as opposed to $3,500 as in the previous example). Next, an order for 100 shares of SBUX is entered into the system. With a volatility of 29.63%, the probability of the order executing is approximately 8.61% resulting in a net change of buying power of -$430.48. In other words, using this approach, the SBUX order is not rejected by the system as it was in the previous example.

While this approach does make adjustments to account buying power as a function of the probability of an order executing, it still does not provide relief for negatively correlated orders. The calculation for the change in buying power due to an open order has no component that accounts for any kind of correlation, negative or otherwise, with any other open orders in a given account.

Alternative Approach Accounting for Dependence In Price Movements

The second alternative approach we explore builds upon the first alternative approach but lifts the assumption that price movements of assets are independent (i.e. uncorrelated). Furthermore, the methodology of this second approach differs completely from that of the first in that the second approach relaxes the assumption that asset returns have normal distributions. Finally, this second approach is significantly more empirical in nature while the methodology of the first alternative approach is more theoretical- the reasons for which will be discussed later in this section.

Again, as in the previous approach, the brokerage determines in advance to what time horizon margin call risk should be managed. Let this time horizon imply a period *P* that will be used to sample historical price data for all assets that are specified in a set of open orders associated with an account. Let O represent this set of n open orders, where x_i is the i-th element, or open order, of the set. Every element x_i has two attributes: an asset a specified by the order and a limit price l at which the order will execute, also specified by the order.

At a moment in time t_o, price values are sampled at a frequency $1/P$ for every unique asset a_j implied by the set O. Assume that prices are sampled backward through time, beginning at t_o and continuing to a time t_ω such that $t_o > t_\omega$ and the volatility of each asset a_j varies similarly to its historical behavior over the time interval defined by t_o and t_ω.

This time interval has $[(t_o - t_\omega) / P]$ sub-intervals t_s of length P. Continuously-compounded returns $r(t_s, a_j)$ are calculated across all sub-intervals for each asset a_j; these returns comprise an empirical joint distribution across all assets implied by the set O and are defined formally as:

$$r(t_s, a_j) = \ln\left(\frac{price(t, a_j)}{price(t - t_s, a_j)}\right)$$

$$price(t,a_j) := price\ of\ a_j\ at\ time\ t$$

Furthermore, let a baseline (continuously-compounded) return b_i be calculated for each order in O, where p_j is the current market price for asset a_j as specified by order x_i:

$$b_i = \ln\left(\frac{l_i}{p_j}\right)$$

For each sub-interval t_s, iterate across each open order x_i in O and evaluate whether the return associated with the sub-interval for the asset a_j associated with x_i is equal to or greater than the baseline return b_i. Let the result of this evaluation $c_i(t_s)$ be defined as follows[3]:

[3] The inequalities used in this evaluation are applicable to buy side orders. The inequalities are reversed for sell side orders.

$$c_i(t_s) = \begin{cases} size_i \times l_i, & r(t_s, a_j) \geq b_i \\ 0, & r(t_s, a_j) < b_i \end{cases}$$

Let $d(t_s)$ be a discrete element of the empirical joint distribution across O and be defined as the sum of evaluations $c_i(t_s)$ for a sub-interval t_s:

$$d(t_s) = \sum_{i=1}^{n} c_i(t_s)$$

Given these definitions, the empirical joint distribution across O will have a maximum of 2^n distinct values that any $d(t_s)$ may take on. Let the set of distinct values that any point in the empirical joint distribution may take on be denoted as E:

$$E_i \leftarrow d(t_s), |E| \leq 2^n$$

The empirical joint distribution will have a series of $[(t_o - t_\omega) / P]$ discrete points, where each point value is assignable to a member of E. In other words, there are up to 2^n distinct combinations of open orders $d(t_s)$ in O possibly executing within an interval of time P from a moment in time t_o, where any $d(t_s)$ maps to an element in E.

Let η_k be the number of times $d(t_s)$ maps to a particular element E_i. The probability of any E_i actually occurring (i.e. the probability of the specific combination of open orders in O implied by E_i being executed within an interval of time P) is then the probability that the buying power of an account should be debited by E_i and can be formally defined as:

$$\Pr(\Delta BP = E_i) = \frac{\eta_k(E_i)}{\frac{(t_o - t_\omega) - P}{P}}$$

Since the point value element E_i implied by $d(t_s) = 0$ is the case where the complete set of orders in O does not execute in the current time period, the joint probabilities for specific combinations of open orders in O executing within an interval of time P are normalized as follows to produce conditional probabilities, where the condition is that at least one open order executes within the interval of time P:

$$\Pr'(\Delta BP = E_i) = \frac{\Pr(\Delta BP = E_i)}{1 - \Pr(\Delta BP = 0)}, E_i \neq 0$$

With these definitions in place, the expected value of the change in buying power over an interval of time P due to a set of open orders O on an account is:

$$E[\Delta BP]\|_P = \sum_{i=1}^{|E|} E_i \times \Pr'(\Delta BP = E_i)$$

Figure 3 lays out an example of this second approach:

Side	Ticker	Size	Limit	Current	Baseline Return
Buy	AAA	100	35.00	37.00	-5.56%
Buy	BBB	100	50.00	54.00	-7.70%
Buy	CCC	100	26.00	29.00	-10.92%

Figure 3

In this example, a customer account has a buying power of $10,000 and submits three open limit orders for stocks AAA, BBB, and CCC, respectively. All orders have a size of 100 shares. The order for AAA has a specified limit price of $35 per share with a current market price of $37 per share. The order for BBB has a specified limit price of $50 per share with a current market price of $54 per share. The order for CCC has a specified limit price of $26 per share with a current market price of $29 per share. This set of open orders O has $n = 2$ elements that specify assets $\{a_1 = AAA; a_2 = BBB, a_3 = CCC\}$. The baseline returns are calculated as:

$$b_1 = \ln(\frac{35}{37}) \cong -5.56\%$$

$$b_2 = \ln(\frac{50}{54}) \cong -7.70\%$$

$$b_3 = \ln(\frac{26}{29}) \cong -10.92\%$$

Assume that the time horizon for managing margin call risk is determined to be $P = 30$ minutes and that prices for the assets AAA, BBB, CCC are sampled at a frequency of 1 / P = (1 / (1 hr / 2)) = twice per hour, or every 30 minutes, for up to t_o - t_ω = 5 hours into the past (for the sake of brevity in this example). Figure 4 shows this data:

t_s	AAA	BBB	CCC	$r(t_s, AAA)$	$r(t_s, BBB)$	$r(t_s, CCC)$
-0.5	38.72	53.94	29.54	-10.48%	-16.99%	-15.00%
-1.0	42.99	63.93	34.32	1.36%	-1.45%	15.65%
-1.5	42.41	64.86	29.35	-2.76%	0.77%	-5.24%
-2.0	43.60	64.37	30.92	3.18%	0.82%	-0.50%
-2.5	42.23	63.84	31.08	8.89%	3.72%	-12.23%
-3.0	38.64	61.51	35.12	-17.74%	11.17%	-5.89%
-3.5	46.14	55.00	37.25	-5.12%	3.29%	-10.90%
-4.0	48.56	53.22	41.54	-1.93%	-1.08%	5.27%
-4.5	49.51	53.80	39.41	2.96%	2.88%	5.32%
-5.0	48.06	52.27	37.37			

Figure 4

In addition to historical prices, Figure 4 shows the empirical joint distribution of continuously-compounded returns $r(t_s, a_j)$ for each of the assets for the sub-intervals over the total historical sampling interval. Each return for each sampling period is compared against the appropriate baseline return b_i to determine values for $c_i(t_s)$. Then, for each sampling period, the set of values for $c_i(t_s)$ are summed to calculate values for $d(t_s)$. These values are displayed in Figure 5:

t_s	$r(t_s, AAA)$	$r(t_s, BBB)$	$r(t_s, CCC)$	$c_1(t_s)$	$c_2(t_s)$	$c_3(t_s)$	$d(t_s)$
-0.5	-10.48%	-16.99%	-15.00%	3,500.00	5,000.00	2,600.00	11,100.00
-1.0	1.36%	-1.45%	15.65%	-	-	-	-
-1.5	-2.76%	0.77%	-5.24%	-	-	-	-
-2.0	3.18%	0.82%	-0.50%	-	-	-	-
-2.5	8.89%	3.72%	-12.23%	-	-	2,600.00	2,600.00
-3.0	-17.74%	11.17%	-5.89%	3,500.00	-	-	3,500.00
-3.5	-5.12%	3.29%	-10.90%	-	-	-	-
-4.0	-1.93%	-1.08%	5.27%	-	-	-	-
-4.5	2.96%	2.88%	5.32%	-	-	-	-
-5.0							

Figure 5

Note than as O has $n = 3$ elements, the empirical distribution of changes to buying power could have had up to 2^n distinct values than any $d(t_s)$ might take on. In other words, in one particular time period, $c_1(t_s)$ might have evaluated to \$3,500 while $c_2(t_s)$ and $c_3(t_s)$ evaluated to zero. Then in another time period, $c_2(t_s)$ might have evaluated to \$5,000 while $c_1(t_s)$ and $c_3(t_s)$ evaluated to zero, and so forth. As $c_i(t_s)$ has a binary evaluation, n open orders results in a maximum magnitude of 2^n for E, the set of distinct values that any point in the empirical joint distribution may take on.

So in this example, E has four elements: {0, 2600, 3500, 11100}. Every $d(t_s)$ is assignable to one of these values. At this point, the probabilities of buying power being debited by an amount equal to any of these values can be calculated as follows:

$$\Pr(\Delta BP = E_1 = 0) = \frac{\eta_k(E_1)}{\frac{(t_o - t_\omega) - P}{P}} = \frac{6}{\frac{4.5}{0.5}} \cong 66.67\%$$

$$\Pr(\Delta BP = E_2 = 2600) = \frac{\eta_k(E_2)}{\frac{(t_o - t_\omega) - P}{P}} = \frac{1}{\frac{4.5}{0.5}} \cong 11.11\%$$

$$\Pr(\Delta BP = E_3 = 3500) = \frac{\eta_k(E_3)}{\frac{(t_o - t_\omega) - P}{P}} = \frac{1}{\frac{4.5}{0.5}} \cong 11.11\%$$

$$\Pr(\Delta BP = E_4 = 11100) = \frac{\eta_k(E_4)}{\frac{(t_o - t_\omega) - P}{P}} = \frac{1}{\frac{4.5}{0.5}} \cong 11.11\%$$

$$100\% = \sum_{i=1}^{4} \Pr(\Delta BP = E_i)$$

Since $E_1 = 0$ is a trivial case, the joint probabilities are normalized to factor out these occurrences to produce the following *conditional* joint probability values:

$$\Pr'(\Delta BP = E_2 = 2600) = 33.33\%$$
$$\Pr'(\Delta BP = E_3 = 3500) = 33.33\%$$
$$\Pr'(\Delta BP = E_4 = 11100) = 33.33\%$$

Finally, the expected value of the change in buying power over a time horizon P can be calculated as:

$$E[\Delta BP]\big|_P = (2600 \text{ x } 33.33\%) + (3500 \text{ x } 33.33\%) + (11100 \text{ x } 33.33\%) = \$5{,}733.33$$

As the customer account has a buying power of \$10,000, the complete set of orders is accepted and buying power is debited by \$5,733.33.

This approach can be further extended to handle situations in which the cancellation of one or more current open orders may be required in order for an account to maintain a positive buying power, where the account's buying power might have 'drifted' to a negative value due to market price movements. Suppose an account has a set of open orders O that has remained fixed (i.e. no members of the set have been removed from the set, no new members have been added to the set) over some time period P. However, also assume that the current market price p_j for at least one asset a_j implied by the set of open orders O has ticked at least once during the time period P. As prices relevant to the set of open orders have changed, the probabilities of these open orders have also changed and therefore, account buying power is impacted.

Let the incremental change in buying power by defined as follows:

$$\Delta BP_I\big|_P = E[\Delta BP]\big|_P - E[\Delta BP]\big|_{P_{-1}}$$
$$E[\Delta BP]\big|_{P_{-1}} \neq 0$$

In other words, the incremental change in buying power over a period *P* is the difference between the expected value of the change in buying power over an interval of time *P* due to a set of open orders *O* and the last non-zero expected value of the change in buying power from a previous time period P_{-1}. If the time period immediately prior to P had an expected change in buying power of zero, then the previous consecutive time period in the past is referenced for a non-zero expected change in buying power and so on until a previous time period is encountered with a non-zero expected change in buying power.

For example: In period P_0, $\Delta BP_{current}$ is -100 while $\Delta BP_{previous}$ is zero; in period P_{-1}, $\Delta BP_{current}$ is 0 and $\Delta BP_{previous}$ is -100; in period P_{-2}, $\Delta BP_{current}$ is -50 and $\Delta BP_{previous}$ is still -100 since in period P_{-1} $\Delta BP_{current}$ was zero.

With an expression for the incremental change in buying power in place, let the buying power after a time period P be calculated as follows:

$$BP(P) = BP(P_{-1}) + \Delta BP_I \big|_P$$

If *BP(P)* is less than zero, then one or more open orders in *O* must be cancelled in order to bring buying power for the account to a positive value. The order in which open orders must be cancelled may be determined by ranking each of the elements in *O* by their respective marginal probability (per the empirical joint distribution) of being executed over the course of a time period *P*. Open orders that are least likely to be executed are cancelled first. After an order in *O* is cancelled, buying power BP is impacted by a new calculated value for the incremental change in buying power as described above. This process is repeated until either buying power is a positive value or the set of open orders *O* is empty.

The primary benefit of this second alternative approach is three-fold: 1) it side-steps the assumption of normality that is inherent in classical risk management theory and thereby implicitly allows for arbitrary distributions, 2) it requires no explicit mechanism to handle variations in asset volatility over time; each constructed empirical distribution incorporates any volatility fluctuations that have occurred in the recent past and assumes similar volatility characteristics over a time horizon *P*, and 3) from an implementation standpoint, this approach is significantly more scalable than an approach that accounts for dependence in price movements between assets using closed-solution statistical (e.g. Pearson) correlations. This approach suggests that a system implementing this solution as part of an automated risk-management system would see a theoretical performance degradation that is linear in nature. In other words, as the number of open orders with unique assets increased, the number of calculations the system would have to repeatedly perform to implement the solution would increase in a linear fashion. However, another approach that, for example, might require utilization of a correlation matrix to manage the dependency measurements between all assets implied by a set of open orders suggests a theoretical performance degradation that would be exponential in nature.

It should be noted that in the special case when the set of open orders O is 1, this approach mimics the current BlueSky Finanical system due to the fact that conditional probabilities are being used. In other words, when an account has only one open order, the expected change to buying power is calculated under the assumption that at least one order executes (i.e. the probability of the order executing is assumed as 100%). As more orders are added the methodology becomes less conservative than the current method.

3. Approaches To Monitoring Margin

Overview of Current Approach

BlueSky Financial does not currently monitor margin levels across its customer accounts as part of its formal risk management strategy. Rather, the firm monitors profit and loss (P&L) numbers and uses those values as the barometer for measuring the financial 'health' of its customer accounts.

However, BlueSky Financial must still adhere to the rules established by the National Associated of Securities Dealers (NASD) and Securities and Exchange Commission (SEC) for margin accounts (see Appendix A for a detailed enumeration of these rules). Further, these margin rules are enforced by the clearing firm that clears all of BlueSky Financial's trades and it is actually the clearing firm that delivers margin calls to customer accounts, as necessary. Currently, the clearing firm has the power to close any customer account that has fallen into a margin call situation at least two times. In fact, immediately upon margin call, the firm has the right to begin liquidating assets to meet margin requirements. However, accounts are usually given a grace period in which to meet margin. A typical example[4] would be for an account in a margin call to have 24 hours to meet the margin requirements before which the account's margin is raised from 25% to 50%; after 5 business days of not meeting the requirements, the account would no longer be allowed to trade on margin at all.

Suggestions for An Alternative Approach

Providing both customers and those managing their accounts (e.g. a risk management desk) up-to-date information on how close an account is to a margin call allows the brokerage to improve customer service and proactively avoid having accounts closed by the clearing firm. To provide this service, BlueSky Financial should constantly monitor a probability of a margin call for each account and trend the ratio of margin available to margin required. In addition, the hedging relationships between assets within each account need to be documented so that when an account is in or near a margin call, a position in the account is not closed that will increase the account's margin requirements. For example, closing an equity position through a sale would increase margin requirements if the equity position were a part of covered call.

The probability of a margin call for a given account during a specified time period P is a function of the set of open orders O associated with the account, the set of positions A associated with the account, and the amount currently borrowed against the account. The probability of a margin call occurring within a time period P can be calculated by expanding the previously-discussed method of calculating the empirical joint distribution of open order executions over the same time period P.

Let the state of an account's margin situation be denoted by a ratio R, calculated as the sum of the amount of margin required for each position and open order in the account divided by the total margin available to the account. Let the amount of margin required for a position or open order be the value v_i of the position or open order times the margin requirement m_i associated with the position or open order. Let the total margin

[4] http://www.investrade.com/

available to the account be the total value of an account's positions/open orders plus cash *c* in the account minus the amount borrowed against the account *d*:

$$R = \frac{\sum_{i=1}^{n} m_i v_i}{\left(\sum_{i=1}^{n} v_i\right) + c - d}$$

If the ratio *R* is less than 1, the account is not in a margin call situation.

Calculation of this ratio over the course of a trading day will provide a metric that measures the percentage of available margin that has been consumed by an account and, by trending this ratio, both customers and the brokerage's risk management desk can readily discern whether an account appears to be moving toward or away from a margin call situation.

The probability of a margin call over a time interval *P* is given by the probability that the ratio *R* is greater than or equal to one. Consider two time horizons, *F* and *B*, where both time horizons are multiples of time interval *P* and *F* is a forward-looking time horizon and *B* is a backward-looking, or historical, time horizon. The historical time horizon *B* is greater than or equal in magnitude to the forward-looking time horizon *F*.

Historical returns $r(t_s,a_j)$ for each asset and open order associated with the account are calculated across time intervals of length *P* for the historical time horizon *B*, similar to the manner in which they are calculated to estimate the probability of a set of open orders executing as described in the previous section of this paper.

Estimated future prices for each asset and open order associated with the account are then calculated for time intervals of length *P* across the forward-looking time horizon *F*. The total number of estimated future prices per asset for any time interval of length *P* (i.e. the number of elements in the set of possible prices for an asset for one of the forward-looking time intervals, where any element in the set if indexed by the variable *q*) in the forward-looking time horizon *F* is:

$$q := \left[1, \left(\frac{B}{P}+1\right)-\left(\frac{F}{P}-1\right)\right]$$

For each time interval P in the forward-looking time horizon, values of R are calculated for each (q-indexed) price in the set of estimated future prices:

$$R(t_{s+kP})\Big|_q = \frac{\sum_{i=1}^{n} m_i v_i(t_{s+kP}) \prod_{j=1}^{k} r_j}{\left(\sum_{i=1}^{n} v_i(t_{s+kP}) \prod_{j=1}^{k} r_j\right) + c(t_{s+kP}) - d(t_{s+kP})}$$

Using these values of R, the probability of a margin call occurring over the forward-looking time horizon F is calculated as:

$$\Pr\Big|_F = \frac{\sum_{q=1}^{\left(\frac{B}{P}+1\right)-\left(\frac{F}{P}-1\right)} \begin{cases} \left[\sum_{j=1}^{\left(\frac{F}{P}\right)} R(t_{s+jP})\right] > 0, 1 \\ else\ 0 \end{cases}}{\left(\frac{B}{P}+1\right)-\left(\frac{F}{P}-1\right)}$$

As an example, assume the margin requirement for equity is 25% and the account owns 1,000 shares of AAA, 1,000 shares of BBB and has an order for CCC at a limit price of $45. Furthermore, assume the account current has $70,000 borrowed against it. If there is an hour left in the day and a decision is made that a two hour look-back time horizon B is appropriate with time intervals of length P = 15 minutes, then the historical prices and returns across each time interval are as follows:

t_v	AAA	BBB	CCC	$r(t_v, AAA)$	$r(t_v, BBB)$	$r(t_v, CCC)$
0	44.01	54.99	44.78	1.0142	1.0044	1.0004
-0.25	43.40	54.75	44.76	1.0137	1.0033	0.9970
-0.50	42.81	54.57	44.89	1.0043	1.0073	1.0073
-0.75	42.63	54.18	44.56	1.0113	1.0044	1.0096
-1.00	42.15	53.94	44.14	1.0132	0.9979	0.9970
-1.25	41.60	54.06	44.27	1.0175	1.0082	0.9970
-1.50	40.89	53.62	44.40	0.9943	1.0111	1.0116
-1.75	41.12	53.03	43.90	1.0141	1.0000	0.9966
-2.00	40.55	53.03	44.05	0.9999	1.0115	0.9956
-2.25	40.55	52.43	44.24			

Figure 6

These prices and returns imply the following sets of estimated future asset prices for the remaining four 15 minutes time intervals across the forward looking period *F*:

q	$AAA(t_{v+0.25})$	$BBB(t_{v+0.25})$	$CCC(t_{v+0.25})$	$AAA(t_{v+0.50})$	$BBB(t_{v+0.50})$	$CCC(t_{v+0.50})$
1	44.51	55.23	45.21	44.71	55.63	45.54
2	44.60	54.88	44.64	45.10	55.12	45.07
3	44.78	55.45	44.64	45.38	55.33	44.51
4	43.76	55.60	45.29	44.53	56.06	45.16
5	44.63	54.99	44.62	44.38	55.60	45.14
6	44.01	55.63	44.58	44.63	55.63	44.43
q	$AAA(t_{v+0.75})$	$BBB(t_{v+0.75})$	$CCC(t_{v+0.75})$	$AAA(t_{v+1.0})$	$BBB(t_{v+1.0})$	$CCC(t_{v+1.0})$
1	45.32	55.82	45.40	45.96	56.07	45.42
2	45.30	55.52	45.40	45.92	55.70	45.26
3	45.89	55.57	44.94	46.09	55.97	45.26
4	45.12	55.94	45.02	45.63	56.18	45.45
5	45.16	56.06	45.00	45.75	55.94	44.87
6	44.38	56.24	44.94	45.15	56.70	44.81

Figure 7

Using these estimated future price sets, we can determine the value of the assets and the amount borrowed against the account for each forward-looking time interval:

q	$\Sigma v_i(t_{s+0.25})$	$d(t_{s+0.25})$	$R(t_{s+0.25})$	$\Sigma v_i(t_{s+0.50})$	$d(t_{s+0.50})$	$R(t_{s+0.50})$
1	99,747	70,000	0.84	100,340	70,000	0.83
2	144,114	114,640	1.22	100,218	70,000	0.83
3	144,874	114,643	1.20	145,213	114,508	1.18
4	99,367	70,000	0.85	100,590	70,000	0.82
5	144,249	114,622	1.22	99,981	70,000	0.83
6	144,221	114,579	1.22	144,686	114,427	1.20
q	$\Sigma v_i(t_{s+0.75})$	$d(t_{s+0.75})$	$R(t_{s+0.75})$	$\Sigma v_i(t_{s+1.0})$	$d(t_{s+1.0})$	$R(t_{s+1.0})$
1	101,139	70,000	0.81	102,028	70,000	0.80
2	100,812	70,000	0.82	101,619	70,000	0.80
3	146,395	114,935	1.16	102,062	70,000	0.80
4	101,059	70,000	0.81	101,814	70,000	0.80
5	101,215	70,000	0.81	146,561	114,869	1.16
6	145,562	114,941	1.19	146,668	114,808	1.15

Figure 8

Using these sets of values for R for each forward-looking time interval, the anticipated instances of margin calls are determined:

$C(t_{s+0.25})$	$C(t_{s+0.50})$	$C(t_{s+0.75})$	$C(t_{s+1.0})$	$C\|_P$
0	0	0	0	0
1	0	0	0	1
1	1	1	1	1
0	0	0	0	0
1	0	0	0	1
1	1	1	1	1

Figure 9

Using the table above, there is a 2/3 chance of a margin call occurring over the remainder of the day and the call will most likely occur in the first fifteen minute forward-looking time interval, $t_{s+0.25}$, if it occurs. This example is by no means calibrated correctly and is only presented as a walk-through of the methodology.

To calibrate the model, the model should be run with subintervals and look-back periods of varying length to benchmark predicted margin calls against actual historical margin calls. Subinterval and look-back length pairs should be indexed against the processing time required to compute the probability calculation:

- The look-back periods should be examined for monotonic asset movements and missing quotes. Look-back periods with monotonic asset movements should be increased in size.

- If predicted margin call probabilities are consistently higher than what are implied by realized margin calls, subinterval lengths should be decreased and vice-versa when predicted margin call probabilities are consistently lower than what are implied by realized margin calls.
- If processing times for the calculation of margin call probabilities are too long, subinterval lengths should be increased.

These model calibrations should be repeated until optimal results are obtained (in terms of the variance between predicted margin calls and realized margin calls, processing time for probability calculations) as defined by a risk management desk.

4. Implementation Considerations

The following sections enumerate the general algorithms which describe how an automated risk management system would 1) manage account buying power, 2) manage historical price data, 3) manage cancellations of open orders, and 4) monitor account margin.

Managing Account Buying Power

- Upon receiving a new order request from a customer account:
  - If the current account buying power (BP) is less than or equal to zero, automatically reject the new order request.
  - Otherwise, evaluate the impact of the new order request on account buying power as follows to determine whether to accept or reject the new order request:
    - Calculate baseline returns for all open orders associated with the account, including the new order request.
      - Baseline return = ln(limit price / current market price)
    - Retrieve 1 year[5] of historical price data from the price database (see the next section on Managing Historical Price Data for a

[5] The use of one year of historical price data to calculate the empirical joint distribution of open orders on customer accounts is arbitrary. The volume of historical price data to use should be a system parameter that

detailed description of the price database and how it maintains current price data); historical price data must be retrieved for each open order currently associated with the account including the new order request that is being evaluated.

- Historical price data will be returned as observations of time intervals of length P, where P is determined by the brokerage as the expected maximum time that would be required for the system to cancel all open orders currently associated with the account[6].
- For each time interval (i.e. observation):
  - Obtain the return minimums and return maximums.
  - For each asset:
    - Compare the historical return associated with the time interval against the baseline return calculated for the asset.
      - Buy-side orders compare against return minimums
      - Sell-side orders compare against return maximums
    - Buy-side orders with a baseline return greater than the time interval's calculated return minimum are considered as executed for the time interval.
    - Sell-side orders with a baseline return less than the time interval's calculated return maximum are considered as executed for the time interval.
  - Calculate the sum of open orders considered as executed for the time interval.

---

can be modified by a risk desk manager as necessary for the system to exhibit optimal risk management characteristics (as determined by the risk desk manager).

[6] The system could be implemented such that P is a system parameter that can be modified by a risk desk manager. Additionally, the system could be implemented such that P is automatically inferred by the system itself based on historical response time measurements taken from previous order cancellations executed by the system.

        - Determine the joint probability distribution across all of the time intervals.
        - Calculate frequencies for each unique open-orders sum calculated across all of the time intervals.
        - Use frequencies to calculate the conditional probabilities of each unique open-order sum occurring (i.e. calculate the probability of a given set of open orders executing in combination within a time period P).
        - Use these conditional probabilities to calculate the impact to buying power (BP). If calculated BP is less than zero, reject the new order request; otherwise, accept the new order request and submit the new order to the order routing system.

Managing Historical Price Data

- An historical price database maintains a revolving set of historical prices for all securities that are tradable through the system.
- For a time period P, real-time price ticks are captured for each asset.
- At the end of the time period P:
    - For each asset:
        - The minimum and maximum price points are determined for the current time period P:
            - MinPrice(P)
            - MaxPrice(P)
        - Calculate the continuously-compounded returns of the minimum price points and maximum price points as follows:
            - $Return_{min}(P) = MinPrice(P) / MinPrice(P_{-1})$
            - $Return_{max}(P) = MaxPrice(P) / MaxPrice(P_{-1})$
        - If no price tick occurred for an asset for the time period P, the minimum and maximum price points from the previous

        - observation are carried over and returns for the current period are considered to be zero.
        - The observation comprised of 1) the current minimum/maximum price points and 2) the calculated return minimum/return maximum is added to the price database. The oldest observation is dropped from the database.

Managing Cancellations of Open Orders

- A process that is responsible for determining whether one or more open orders associated with an account should be cancelled is triggered upon one of the following events:
    - A period of time P passes in which at least one asset specified by any open order associated with the account has experienced a net change in price since the previous period.
    - A new order request is submitted to the system.
- This process is responsible for maintaining a short history of changes in buying power (ΔBP) per account, where ‘short history’ is defined as the last non-zero change in buying power.
- For the account on which the process was triggered:
    - Utilize the algorithm for managing account buying power to calculate a change in buying power ΔBP’.
    - If ΔBP’ is a non-zero value:
        - Calculate the incremental change in buying power by taking the difference between ΔBP’ and the last non-zero change in buying power (ΔBP) maintained in the process’s history:
            - Incremental ΔBP = ΔBP’ – ΔBP
        - Add the incremental ΔBP value calculated to the account’s buying power.
        - DO: If the new value of buying power for the account is less than zero:

- If there are no open orders associated with the account, exit the DO-LOOP. Otherwise, identify the open order associated with the account least likely to execute, where open orders are ranked within the account by their empirical marginal probability of being executed over a period of time P.
- Cancel the least likely open order.
- Recalculate the incremental change in buying power with one less open order associated with the account.
- Add the incremental ΔBP to the account's buying power.
- LOOP

Monitoring Account Margin

- Determine the time interval you wish to measure the probability, T, the subinterval that you would like to track changes in value and executions, t, and the look-back period you want to use to calculate the probability, H.
- Using a dataset of historic returns similar to the Order Flow determine the possible paths the price could take till the end of the desired time interval.
- For each time interval, t_i, and position/order, determine value by multiplying current price*historic return
- Determine the new margin requirements for the asset value and executed orders
- If margin requirement is greater than margin available, then label a margin call.
- Recursively repeat the steps above compounding returns:
    - 2nd Period Prices=Prices Period 1*Historic Returns Period 2, 3rd Period Prices =2 Period Prices * Historic Returns Period 3 … until you reach the number of time periods until the end of the day – T/t.
- If the record has any margin calls across any time period label it a margin call.
- Get the probability of a margin call by taking total margin calls/total time intervals, where the total time intervals will be H/T.
- The expected time to a margin call can also be estimated using this sequence.

- Loop through the methodology every time there is a price tick or as calibration determines necessary.

5. Areas For Improvement

It must be noted that the approach outlined in this paper should only be considered as a first iteration of incremental improvements to BlueSky Financial's current system. This approach could be further refined, for example, by incorporating a measure of the cost to the firm of shutting down accounts due to margin calls vs. the potential increase in revenues from trading commissions.

In that vein, the following are some additional areas in which the model proposed in this paper could be further developed in subsequent iterations:

- **Benchmarking**. The models presented for calculating impact to buying power and probability of an account entering margin call should be calibrated using historical trade information from the brokerage to determine optimal time period durations to use in these calculations.
- **Risk Tolerance/VAR**. The current approach also assumes that all customers and the firm have the same risk tolerance. In order, to provide a more robust solution the risk tolerance of the brokerage and their customers should be considered and a methodology such as value at risk evaluated.
- **Customer Preference**. There is a human element to trading and customer feedback should be gathered when first implementing any new rules or methods.
- **Adjusting for Liquidity**. The calculations used to measure the buying power impact of open orders and the calculations used to monitor accounts for potential margin calls situations could be refined to explicitly adjust for illiquid assets. Once a methodology for identifying illiquidity is determined, the approach for impacting buying power should be modified such that orders involving illiquid assets impact buying power

conservatively, or in other words, assuming that the probability of their execution is 100% (i.e. current approach).

- **Adjusting for Selection Bias**. The current models for margin and open order management include a selection bias due to the fact that only assets viable at the time of calculation are considered. There are no adjustments made for delistings or bankruptcies. To make these models more robust, an adjustment should be made based upon the probability of an asset jumping to a value of zero. This adjustment could be calculated by using all assets of similar Fama-French[7] factors and determining the historic rate of bankruptcies. Calculated order fill probabilities or possible asset values would by adjusted accordingly.
- **Bayesian Estimation**. The current solution of using historic data may become problematic for an institutional customer or very large account. If Bayesian Markov Chain Monte Carlo[8] methods are used to create a joint distribution using software such as WinBUGS[9], a faster calculating, less data intensive solution could be implemented.

[7] Fama French factors provide industry and size classifications

[8] Markov Chain Monte Carlo methods have been studied as alternatives to classical statistics and are often used in problems where complex functions, curve fitting, and integration are required.

[9] WinBUGS is a shareware Bayesian statistics program available from the Imperial College School of Medicine

Appendix A

Basic Requirements for Margin Accounts

NASD/SEC have defined the following basic requirements that margin accounts must adhere to in order to avoid margin call situations:

- Trading accounts that are categorized as day-trader[10] accounts must maintain a minimum equity level of $25,000. Otherwise, trading accounts must maintain a minimum equity level of $2,000.
- Trading accounts must maintain at least 25% of the current value of their long positions (excluding futures securities)
- For short positions in stocks that have a market value of $5/share or less, trading accounts must maintain either $2.50/share or 100% of the current market value of the short position, whichever is greater.
- For short positions in stocks that have a market value of $5/share or greater, trading accounts must maintain either $5/share or 30% of the current market value of the short position, whichever is greater.
- For short positions in bonds, trading accounts must maintain either 5% of the principal or 30% of the current market value of the short position, whichever is greater.
- Trading accounts must maintain 20% of the current market value of all long and short futures contracts.

In addition, there are a set of rules that are applicable to short positions on options of various underlying types:

- For options on equities and industry indexes, trading accounts must maintain 20% of the value of the equities or industry indexes.

[10] Trading accounts are typically classified as day-trading accounts when four or more *day trades* are executed through the account within five business days. Day trades are defined as the purchase and sale of the same position within a single trading day. Day trades must comprise at least 6% of an account's trades for a given trading day. Brokerages have the flexibility to classify an account as a day-trading account even if these conditions are not necessarily met.

- For options on broad indexes, trading accounts must maintain 30% of the value of the broad indexes.
- For options on T-Bills, trading accounts must maintain 0.35% of the value of the T-Bills.
- For options on T-Notes, trading accounts must maintain 3% of the value of the T-Notes.
- For options on T-Bonds, trading accounts must maintain 3.5% of the value of the T-Bonds.
- For options on foreign currencies, trading accounts must maintain 4% of the value of the foreign currency.
- For options on interest rate contracts, trading accounts must maintain 10% of the value of the contracts.

Margin requirements may be reduced for options that are out-of-the-money.

Complete details regarding margin requirements as specified by NASD can be referenced at the following URL:

`http://nasd.complinet.com/nasd/display/display.html?rbid=1189&record_id=1159000525&highlight=margin+rule`

Appendix B

Three Simulations Using the Empirical Joint Distribution Approach

Two sets of simulations were run to test the applicability of the empirical joint distribution approach for managing open orders using the stocks listed below in Figure 10:

PFE	Pfizer
C	Citicorp
DELL	Dell
CD	Cendant
MSFT	Microsoft

Figure 10

Simulation set 1 considered a combination of these stocks (short or long) specifically constructed to produce low probabilities of executions whereas simulation set 2 considered a combination of these stocks constructed to produce high probabilities of execution. Market data from December 6, 2004 was assumed as present-time information. A look-back period of 11 months was used and divided into 10 minute subintervals. Weekends, holidays, and non-trading hours were adjusted for by assuming the time between any two consecutive trading sessions was a 10 minute subinterval. Gaps in returns data for any subinterval assumed a zero-percent return from the preceding subinterval.

The results of the examples are as expected. The low probability orders have less buying power impact and those negatively correlated receive relief. In addition, the results seem to indicate that as more assets are added, the conditional probability approaches the expected value calculation. If order portfolios are large enough, this may provide potential relief for the calculation algorithms. Another finding is that as the probability of execution decreases, the short-term order executions act as if they are independent. This provides another possible way to save calculation time by assuming that extremely low probability orders are independent.

Low Probability Simulations: Transaction Summary

Portfolio 1.1						
Action	Ticker	Size	Limit	Current Price	BP Impact	Return
Buy	PFE	1000	$ 26.00	$ 27.25	$ 26,000.00	0.95412844
Short Sell	C	1000	$ 46.00	$ 45.70	$ 23,000.00	1.006564551
Total					$ 49,000.00	
New Method					$ 23,187.50	
Portfolio 1.2						
Action	Ticker	Size	Limit	Current Price	BP Impact	Return
Buy	DELL	500	$ 40.00	$ 41.55	$ 20,000.00	0.962695548
Buy	C	500	$ 44.50	$ 45.70	$ 22,250.00	0.973741794
Total					$ 42,250.00	
New Method					$ 22,826.09	
Example 1.3						
Action	Ticker	Size	Limit	Current Price	BP Impact	Return
Buy	DELL	500	$ 40.00	$ 41.55	$ 20,000.00	0.962695548
Buy	CD	500	$ 21.00	$ 22.80	$ 10,500.00	0.921052632
Total					$ 30,500.00	
New Method					$ 20,200.00	

Figure 11

High Probability Simulations: Transaction Summary

Portfolio 2.1						
Action	Ticker	Size	Limit	Current Price	BP Impact	
Buy	PFE	1000	$ 27.20	$ 27.25	$ 27,200.00	0.998165138
Short Sell	C	1000	$ 45.50	$ 45.70	$ 22,750.00	0.995623632
Total					$ 49,950.00	
New Method					$ 31,518.87	
Portfolio 2.2						
Action	Ticker	Size	Limit	Current Price	BP Impact	Return
Buy	DELL	500	$ 41.50	$ 41.55	$ 20,750.00	0.998796631
Buy	C	500	$ 45.74	$ 45.70	$ 22,870.00	1.000875274
Total					$ 43,620.00	
New Method					$ 17,589.22	
Portfolio 2.3						
Action	Ticker	Size	Limit	Current Price	BP Impact	Return
Buy	DELL	500	$ 40.00	$ 41.55	$ 20,000.00	0.962695548
Buy	CD	1000	$ 22.70	$ 22.80	$ 22,700.00	0.995614035
Total					$ 42,700.00	
New Method					$ 22,020.19	

Figure 12

References

1. Robert L. McDonald, *Derivates Markets*, Addison Wesley 2003.
2. Steve L. Allen, *Financial Risk Management: A Practitioner's Guide to Managing Market and Credit Risk*, Wiley February 14, 2003.
3. Sergio M. Focardi & Frank J. Fabozzi, *The Mathematics of Financial Modeling and Investment Management*, Wiley March 19, 2004.
4. Keith E. Muller& Bethel A. Fetterman, *Regression and ANOVA: An Integrated Approach Using the SAS System*, Wiley-SAS May 23, 2003.
5. Donald R. Van Deventer, Kenji Imai, Mark Mesler, *Advanced Financial Risk Management: Tools & Techniques for Integrated Credit Risk and Interest Rate Risk Managements*, John Wiley & Sons November 19, 2004.
6. Torben G. Andersen, Tim Bollerslev, Peter F. Christoffersen and Francis X. Diebold, *Practical Volatility and Correlation Modeling for Financial Risk Management*, PIER Working Paper 05-007 January 11, 2005.
7. David F. Babbel, Anthony M. Santomero, *Risk Management by Insurers: An Analysis of the Process*, Wharton Financial Institutions Center May 1996.
8. *Managing Market Risk in Banks*, Reserve Bank of Australia Bulletin December 1996.