РАЗДЕЛ 1
НАУЧНЫЕ ИССЛЕДОВАНИЯ И КОНЦЕПЦИИ

SECTION 1
ACADEMIC INVESTIGATIONS & CONCEPTS

ASYMMETRIC INFORMATION, TRADING VOLUME, AND PORTFOLIO PERFORMANCE

Antony Jackson*

Abstract

In dealership markets, asymmetric information feeds through to higher transaction costs as dealers adjust their bid-ask spreads to compensate for anticipated losses. In this paper, we show that the presence of asymmetric information can also provide a positive externality to those market participants who operate in multiple markets-portfolio managers. Specifically, insiders lower the estimation errors of portfolio selection methods, thus improving asset allocation. We develop multiple artificial markets, in which portfolio managers trade alongside informed and uniformed speculators, and we contrast the performance of 'volatility timing'—a method that relies on efficient price discovery - with that of 'naive diversification'. Volatility timing is shown to consistently outperform naive diversification on a risk-adjusted basis.

Keywords: Asymmetric Information, Portfolio Selection, Stochastic Simulation

JEL Classification: D82, G11, G12

* *School of Economics, University of East Anglia, Norwich NR4 7TJ, United Kingdom*
Tel.: +44 (0)1603 59 3876
E-mail: antony.jackson@uea.ac.uk

1. Introduction

In market microstructure models, transaction costs arise endogenously - either through the inventory management process of the monopolist (Ho and Stoll, 1981), or through the asymmetric information advantage of insiders (Glosten and Milgrom, 1985). Repeated iteration of the Glosten and Milgrom (1985) model generates intra-day price dynamics via the price setting behavior of a market maker responding to the flow of orders arriving from a large pool of informed and uninformed traders. The degree to which intra-day prices 'discover' true fundamental value depends on how sensitive the dealer's priors are to the flow of new orders. The dealer adjusts prices most rapidly when the proportion of informed trade and the volume of orders are high.

In this paper, we study the effects of asymmetric information in the wider context of multiple asset markets. In an individual market, a higher probability of informed trade unambiguously leads to higher transaction costs. We suggest, however, that there are subtle benefits of asymmetric information that accrue to those who operate across many markets: portfolio managers. The reason is that portfolio selection methods rely to various degrees on efficient price discovery - the ability of the market mechanism to accurately reflect underlying fundamentals. We argue that private information counteracts the impediment to price discovery inherent in low trading volume, and that there appears to be an optimal level of private information, given the other characteristics of a particular market.

Our approach is to simulate multiple assets with correlated fundamentals. In our dealership markets, insiders act as the conduit between fundamentals and prices. To assess the costs and benefits of asymmetric information to portfolio managers, we contrast the performance of a strategy that relies on efficient price discovery - the volatility timing strategy - with naive diversification. This choice partly reflects recent developments in the portfolio choice literature, but also reflects our preference for methods that offer the practical advantage of rapid computation.

Attention has recently focused on portfolio management strategies that avoid the problems associated with full mean-variance optimization: singularity in the covariance matrix of returns, and excessive volatility in asset allocations. Restrictions are placed on elements of the covariance matrix, or 'shrinkage' estimators are formed as weighted sums of the sample covariance matrix and a simple 'target' matrix; see, for example, Jagannathan and Ma (2003) and Tu and Zhou (2011). The naive diversification strategy (DeMiguel et al., 2009) entirely removes the need for an estimated covariance matrix, instead allocating an equal share of capital to all portfolio constituents. The volatility timing strategy (Kirby and Ostdiek, 2012) is more involved - basing its allocation on relative volatilities calculated using moving windows of asset prices. Both strategies share the characteristics of full capital allocation and no-short-sales.

The approach of this paper is to take advantage of the simple Bayesian updating mechanism offered by the binomial branching structure of the sequential trade model (Glosten and Milgrom, 1985), while retaining the original statistical properties of the full multivariate simulation of underlying fundamental values. This is achieved by mapping multivariate normal returns into their Bernoulli equivalents, a process that requires boosting the elements of the original covariance matrix (Einrich and Piedmonte, 1991). Our simulation methodology does not place any restrictions on the number of portfolio constituents. Covariance matrices are randomly generated using a wide range of parameter values within a single-index factor model. We generate multivariate asset returns using the Cholesky factorization of these matrices, which requires matrix inversion, but we address the potential singularity problem by reconstructing those matrices with negative eigenvalues (Rebonato and Jackel, 1999).

A further innovation of this paper is to borrow the recombining tree structure of the Cox et al. (1979) binomial options pricing model. We replace the risk-neutral probabilities of Cox et al. (1979) with the probabilities implied by a single-index model with drift. Multiple markets are linked together by the correlations between their fundamental values. The recombining tree structure lays the foundation for future research on the stochastic arrival of information, as it keeps the dealer's Bayesian updating task manageable. Information arrives at the beginning of each trading period, with true values revealed at the end of each period.

We draw an important distinction between the trading population that generates prices (uninformed and informed speculators), and portfolio managers who act upon multiple asset prices. A feature of the Glosten and Milgrom (1985) model is that as the dealer processes orders, the uncertainty of the true underlying value diminishes, in turn leading to narrower bid-ask spreads. If we were to posit portfolio managers as arriving randomly during the session - like the rest of the population - we would also randomly vary the impact of transaction costs. We prefer instead to place all portfolio manager trades at the opening bid- ask spreads of each period, which enables transaction costs to be a pure function of the probability of informed trade. This abstraction also enables us to sidestep the tricky issue of strategic behavior when market participants trade more than a single unit. Portfolio managers in our model are able to accurately signal to the dealer that they are uninformed. In concurrent research, we consider the liquidity cost that must be borne by portfolio managers who are unable to naturally differentiate themselves from the rest of the population. In this version, portfolio managers operate in multiple 'Kyle' auction markets (Kyle, 1985).

The final bid-ask spread of each session is used to calculate the session 'close'. Portfolio managers mark their holdings to market using closing prices. The day-to-day changes in account value imply a series of strategy returns, with mean returns and risk-adjusted returns (Sharpe ratio) following. In addition, volatility timing managers use closing prices in the volatility calculations that determine their asset allocations. This is why trading volume and the probability of informed trade have a joint influence on the performance of the volatility timing strategy. A large flow of orders makes it easier for the dealer's posterior probabilities to converge to the true probabilities, but unless there is a sufficient level of informed trade, even high volume may be insufficient for efficient price discovery. In the extreme, with an entirely uninformed population, a competitive, risk-neutral dealer quotes a single bid/ask price, and sees no reason to adjust the price in response to trading volume. Instead, the price jumps each time the changes in fundamental value become common knowledge.

The model of fundamentals presented in Section 2 generates multivariate normal returns using a single index factor model. Individual assets are characterized by the sensitivity of their returns to movements in the market index, and through the portfolio's correlation matrix. These data determine the sizes and probabilities of 'up' and 'down' movements in our Cox et al. (1979) discretization scheme. Intra-day trade takes place in individual competitive markets that are indirectly connected by the insiders who make decisions based on private access to fundamental information. The latest change in fundamental value is made common knowledge at the end of each day, with dealers adjusting their opening spreads accordingly. Although beyond the scope of the current paper, the recombining structure of the Cox et al. (1979) scheme allows the revelation of information to occur stochastically, whilst keeping the dealer's updating task manageable. A natural way to do this is to use a geometric distribution to randomly select the release of 'news announcements'.

Section 2 also describes the Einrich and Piedmonte (1991) procedure for transforming multivariate normal random variables into their Bernoulli equivalents. We describe the Rebonato and Jackel (1999) method for dealing with singular correlation matrices, and list the parameter assumptions used in constructing our various portfolios.

Section 3 describes the model we use to create intra-day price dynamics and closing prices. We derive probability updating equations in terms of the probabilities of informed trade and the probability of value rising. The sizes of price movements, and their probabilities of occurrence, feed from Section 2.

Once the time series of opening and closing prices has been generated, we test the performance of the naive diversification and volatility timing strategies. In Section 4, upon observing the vector of opening bid-ask quotes, each manager re-values his current positions, and calculates his desired holdings. The naive diversification manager allocates capital equally between assets, whereas the volatility timing manager allocates capital using rolling estimates of volatility.

In Section 5, we present the results, and we use nonparametric methods to identify the key drivers of portfolio performance. The key driver of mean returns is the probability of informed trade, while the key driver of the Sharpe ratio statistic is the strategy type.

The determinants of the highest mean return are intuitively straightforward: substantial volume in illiquid states, combined with low probabilities of informed trade. The determinants of a strategy's Sharpe ratio offer a more interesting story. The Sharpe ratios of the volatility timing strategy dominate those of the naive diversification strategy across all market conditions. Since mean returns are not driven by strategy type, it must be that the volatility timing strategy offers improved risk-adjusted returns via lower risks. There are substantial improvements in the volatility timing strategy's risk-adjusted performance as the number of assets in the portfolio is increased, but the most intriguing driver is the probability of informed trade—the Sharpe ratios corresponding to a 1% probability of informed trade are lower than those corresponding to higher probabilities. Evidently, the volatility timing strategy benefits from the improved price discovery offered by 'reasonable' levels of asymmetric information, but these gains are eventually overwhelmed by higher transaction costs.

The paper concludes with suggestions for future research. In particular, our recombining tree structure allows for staggered news arrivals, without the need for great complexity in the dealer's Bayesian updating problem. The use of a geometric distribution for the timing of news arrivals would seem a sensible start, with insiders maintaining their informational advantage at all times.

2. Fundamentals

The log-returns of the portfolio constituents' fundamental values are multivariate normally distributed. The returns generating process is assumed to be a single-index model, where the return on the risk-free asset is normalized to zero. An individual asset's expected returns are a simple function of its beta coefficient and the expected return to the market index:

$$r_i = \beta_i r_m ,$$

where r_i denotes the expected return to asset i, and r_m denotes the expected return to the market index. The beta coefficient β_i is defined by

$$\beta_i = \frac{\sigma_{i,m}}{\sigma_m^2},$$

and measures the ratio of the covariance of the returns to an asset and those of the market index to the variance of the returns to the market index.

The expected return to the market index is assumed to be constant, $r_m = 10\%$ p.a., with a constant annual volatility of $\sigma_m = 20\%$ p.a. Individual volatilities σ_i , betas β_i , and pairwise correlations $\rho_{i,j}$ are drawn independently from various uniform distributions. Table 1 lists the various specifications. Each asset's annual volatility is assumed to lie in the range 5% to 40%, and its beta coefficient in the range 0.50 to 1.50. The pairwise correlation coefficient between assets lies in the range 0.00 to 1.00. These parameter distributions are chosen to allow for a wide range of volatilities, as well as a variety of relationships with the market index.

Table 1. Simulation Parameter Distributions

Parameter	Description	Value
n	Number of portfolios	1000
N	Number of portfolio constituents	{2, 5, 10}
σ_m	Market Index volatility (p.a.)	20%
r_m	Market Index expected return (p.a.)	10%
σ_i	Asset i volatility (p.a.)	Uniform (5%, 40%)
β_i	Asset i beta	Uniform (0.50, 1.50)
$\rho_{i,j}$	Correlation (i, j)	Uniform (0.00, 1.00)

Each portfolio consists of 2, 5, or 10 stocks. For each of these different portfolio sizes, we simulate 1000 portfolios using randomly-generated correlation matrices. We assume that each $\rho_{i,j}$ $(i \neq j)$ is drawn independently from a continuous uniform distribution with range

[0, 1]. The elements along the main diagonal are set to 1, and those below the main diagonal are set (by symmetry) according to $\rho_{j,i} = \rho_{i,j}$. The resulting correlation matrix **C** is used to generate multivariate normal random variables. In order to be compatible with the simple intra-day sequential trade model, these multivariate random variables are then transformed to Bernoulli random variables.

The sizes of fundamental value movements are described by the following equations:

$$u = \exp\left(\sigma\sqrt{\Delta t}\right) \qquad d = 1/u \tag{1}$$

where σ denotes annual volatility, and $\Delta t \equiv 1/250$ denotes a single day in which prices can move up u or down d, where the size of the down move is simply the reciprocal of the up move.

The probabilities of the moves are calculated using a modified version of the Cox et al. (1979) discretization scheme, in which the risk-neutral drift rate is replaced by the stock's expected return:

$$\text{Prob}(u_i) = \frac{\exp(r_i/250) - d_i}{u_i - d_i} \tag{1}$$

This enables the design of a procedure that starts by generating correlated multivariate random variables, and then maps those variables into a simpler Bernoulli distribution. The binomial process for fundamental value fits comfortably with the sequential trade model of Section 3, which—when iterated over many time periods—recaptures the statistical properties of the original distribution.

Multivariate Bernoulli Transformation

The square matrix **C** can be expressed in terms of its diagonal eigenvalue matrix **Λ**, and the corresponding unit-length eigenvector matrix **S**:

$$CS = S\Lambda \tag{3}$$

Provided the matrix **C** has only non-negative eigenvalues, Equation 3 can be post-multiplied throughout by the inverse matrix $\mathbf{S}^{-1}$ to yield

$$C = S\Lambda S^{-1} \tag{4}$$

Furthermore, since the eigenvector matrix has been defined in terms of unit-length vectors, Equation 4 may be written as

$$C = S\Lambda S^{T} \tag{5}$$

with $\mathbf{S}^{\mathbf{T}}$ replacing $\mathbf{S}^{-1}$. Now define $\mathbf{B} = \mathbf{S}\sqrt{\mathbf{\Lambda}}$. Then Equation 5 may be rewritten as

$$C = S\sqrt{\Lambda}\sqrt{\Lambda}S^{T} = BB^{T}, \tag{6}$$

the spectral decomposition of the correlation matrix. A matrix of correlated standard normal random variables **X** is constructed using the transformation

$$X = BZ \tag{7}$$

where **Z** is a matrix of independent standard normal random variables.

Our objective is to use a simple mapping from the matrix **X** of correlated normal random variables into a matrix **P** of correlated Bernoulli random variables, which in turn are used in the binomial branching structure of the Glosten and Milgrom (1985) sequential trade model. We denote the multivariate Bernoulli distribution's marginal probabilities by $\pi_1, \pi_2, \ldots, \pi_n$. These probabilities correspond to each asset's probability of an up move, as defined by Equation 2. If a stock's characteristics are such that it has a high expected rate of return, then its probability of an up move will be higher - the magnitude of the move is given by Equation 1.

The resulting correlation matrix of returns has pairwise correlation coefficients that are significantly lower than the original correlation matrix **C**. The problem is overcome by first increasing the off-diagonal elements of **C** using the procedure proposed in Einrich and Piedmonte (1991). First, the quantiles of the standard normal distribution are evaluated at the Bernoulli marginal probabilities:

$$z_{\pi_i} = z(\pi_i).$$

Then, the pairwise correlation coefficients of **C** above the main diagonal are replaced by numerically solving for $u_{i,j}$ in the following equation:

$$\Phi_2\left(z_{\pi_i}, z_{\pi_j}; r_{i,j}\right) = \rho_{i,j}\sqrt{\pi_i(1-\pi_i)\pi_j(1-\pi_j)}\;(1-\pi_j) + \pi_i\pi_j, \quad (8)$$

where $\phi_2(\cdot)$ is the c.d.f. of the bivariate standard normal distribution. The correlation coefficients below the main diagonal are set as $\rho_{j,i} = \rho_{i,j}$, ensuring that the new 'boosted' correlation matrix $\mathbf{C}'$ is square-symmetric.

Using Equations 3 through 6, spectral decomposition is performed on **C**. However, it is well known (especially for larger portfolios) that the correlation matrix is likely to have at least one negative eigenvalue, making it impossible to invert the correlation matrix in the first step of the decomposition. One method of addressing this problem is to follow Rebonato and Jackel (1999) in setting any negative eigenvalues to zero, and then reconstructing a new correlation matrix as an approximation to the original. The eigenvector matrix **S** is post multiplied by the square-root of the corrected eigenvalue matrix $\mathbf{\Lambda}'$ to yield the adjusted factor matrix

$$B = \sqrt{T}S\sqrt{\Lambda'},$$

where **T** is a diagonal scaling matrix with elements $t_i = \sum_{j=1}^{n} s_{i,j}\lambda_j$, i.e., the row-wise eigenvectors multiplied by the adjusted eigenvalues. The adjusted correlation matrix $\mathbf{C}''$ is defined by

$$C'' = BB^T.$$

Finally, the boosted matrix of correlated standard normal random variables $\mathbf{X}' = \mathbf{B}'\mathbf{Z}$ is mapped into a matrix of correlated Bernoulli random variables **P** using the rule

$$p_{ii} = 1 \quad \text{if} \quad x_{ii} \le z_{\pi_i}$$
$$= 0 \quad \text{otherwise.}$$

To summarize: we randomly create a target correlation matrix **C** that describes the original multivariate distribution of fundamental returns. The pairwise correlation coefficients of **C** are boosted in order to construct a new matrix $\mathbf{C}'$ to be used in the generation of multivariate Bernoulli random variables. If the eigenvalues of $\mathbf{C}'$ are all non-negative, then spectral decomposition is performed on $\mathbf{C}'$; otherwise, a new correlation matrix $\mathbf{C}''$ is constructed from the 'corrected' diagonal matrix of eigenvalues. The adjusted matrix of correlated standard normal random variables is then mapped into a matrix of correlated Bernoulli random variables, which when used in conjunction with Equations 1 and 2 recovers the properties of the original correlation matrix **C**.

3. Intra-day Trading

The intra-day model is based on Glosten and Milgrom (1985)[1], and is used to generate time series of opening and closing prices, with a view to testing various portfolio strategies. Opening prices are used to revalue current positions, and to determine the prices at which fresh purchases and sales are transacted; the opening bid-offer spread determines transaction costs. Closing prices are prices at which it is not possible to trade, but are commonly the ones used to calculate the returns to a strategy. They also play a central role in the volatility timing strategy, as the strategy uses volatility estimates calculated from rolling windows of closing prices. Closing prices are determined by the set of dealer quotes after the final trade of the day. The price discovery mechanism is expected to function better in high-volume conditions, with aggregate order imbalances reflecting asymmetric information.

There are four market participants: informed traders, uninformed traders, portfolio managers, and risk-neutral dealers. Price competition between dealers ensures that each dealer exactly offsets the expected losses from trading with informed traders with the expected gains from trading with uninformed traders. Provided the details of individual trades are made available to all dealers, the problem reduces analytically to that of one dealer.

Trading volume λ determines the number of trades that take place each day. The sequential trade model deals with daily trading volume as a sequence of single-unit transactions between individual traders and the dealer. Traders are randomly selected, one at a time, from a large pool

[1] Other references include Easley and O'Hara (1992), who extend the model to include the possibility of infrequent information asymmetry, and Back and Baruch (2004).

of informed and uninformed traders, with q denoting the probability of drawing an informed trader. The dealer quotes an ask price at which traders may buy a single unit of the asset, and a bid price at which they may sell. When presented with these quotes, traders have the option to buy, sell, or pass on the trading opportunity. The dealer knows that informed traders will choose to buy only if Ask < V_t (ask is below fundamental value), and will choose to sell only if Bid > V_t (bid is above fundamental value). Uninformed traders choose to trade for reasons unrelated to private information. They are, for example, motivated by hedging requirements, or by the need to meet liabilities. We assume that, for all quotes, uninformed traders randomly buy or sell with probability 1/2.

In the basic version of the model, the dealer learns the true value of the asset at the end of each trading period. In the meantime, his ability to keep track of value depends on liquidity (the number of trades each period), and the proportion q of informed traders.

Figure 1 illustrates the unconditional probabilities of various events, each organized by the trading decisions of informed and uninformed traders - and two possible changes in value.

We assume that the dealer is fully conversant with the structure of the model, and that his specialist knowledge ensures that he uses correct values for volatility and expected returns.

As a consequence, he correctly calculates the unconditional probabilities p and $1-p$ of up and down moves. Informed traders never pass, because the presence of uniformed traders $((1-q) > 0)$ ensures that if the next trader buys, expected value must lie below the 'up' value V^1. This is because the buy trade could come from an uninformed trader in the 'down' value state of the world. Similarly, if the next trader sells, expected value must lie above the 'down' value V^2.

The asset price is initially set to fundamental value V^0, and the returns generating process determines whether value moves up to V^1 or down to V^2. A trader is chosen at random from the pool of informed and uninformed traders, with q denoting the probability of selecting an informed trader, and $1 - q$ the probability of selecting an uninformed trader. Informed traders immediately receive a signal of the new value. The dealer's risk-neutrality, and the zero-profit condition, leads the dealer to set his quotes according to

$$\text{Ask} = \text{E}[V|\text{next trader buys} \qquad (9)$$

And

$$\text{Bid} = \text{E}[V|\text{next trader sells}]. \qquad (10)$$

The ask is set such that the dealer expects to make zero profit if the next trade is a buy. Because of the presence of uninformed traders, buy trades can occur for both values of V. Equation 9 therefore expands to

$$\text{Ask} = V^1\text{Prob}\ V^1\ \text{buy}\ +\ V^2\text{Prob}\{V^2|\text{buy}\} \qquad (11)$$

where $\text{Prob}\ V^1\ \text{buy}$ and $\text{Prob}\{V^2|\text{buy}\}$ serve as the dealer's updating equations in a dynamic setting. Using Bayes' Rule, we obtain

$$\text{Prob}\ V^1\ \text{buy}\ = \frac{\text{Prob}\ \text{buy}\ V^1\ \text{Prob}\ V^1}{\text{Prob}\{\text{buy}\}} \qquad (12)$$

And

$$\text{Prob}\ V^2\ \text{buy}\ = \frac{\text{Prob}\ \text{buy}\ V^2\ \text{Prob}\ V^2}{\text{Prob}\{\text{buy}\}}. \qquad (13)$$

Using the probabilities in the rightmost column of Figure 1,

$$\text{Prob}\ V^1\ \text{buy}\ = \frac{\pi_1(1+q)}{\pi_1\ \ 1+q\ \ +\pi_2\ \ 1-q} \qquad (14)$$

And

$$\text{Prob}\ V^2\ \text{buy}\ = \frac{\pi_2(1-q)}{\pi_1\ \ 1+q\ \ +\pi_2\ \ 1-q} \qquad (15)$$

The bid price is derived in a similar manner. Using the probabilities in the rightmost column of Figure 1,

$$\text{Prob}\ V^1\ \text{sell}\ = \frac{\pi_1(1-q)}{\pi_1\ \ 1-q\ \ +\pi_2\ \ 1+q}$$

and

$$\text{Prob}\ V^2\ \text{sell}\ = \frac{\pi_2(1+q)}{\pi_1\ \ 1-q\ \ +\pi_2\ \ 1+q}$$

Trading volume determines the quantity of random draws from the trading population in each session. The dealer updates his bid-ask spread using the updating Equations 9 and

10. After the final trade of the day, the mid-price of the bid and ask prices is used as the closing price of the day. With a small probability $\zeta = 0.1$, trading volume may change from a low-volume regime to a high-volume regime, and *vice versa*. The fundamental value becomes common knowledge in the time between the close of the current session and the open of the next. The dealer's opening spread reflects this update in public information.

Figure 1. Probability of different types in the sequential trade model

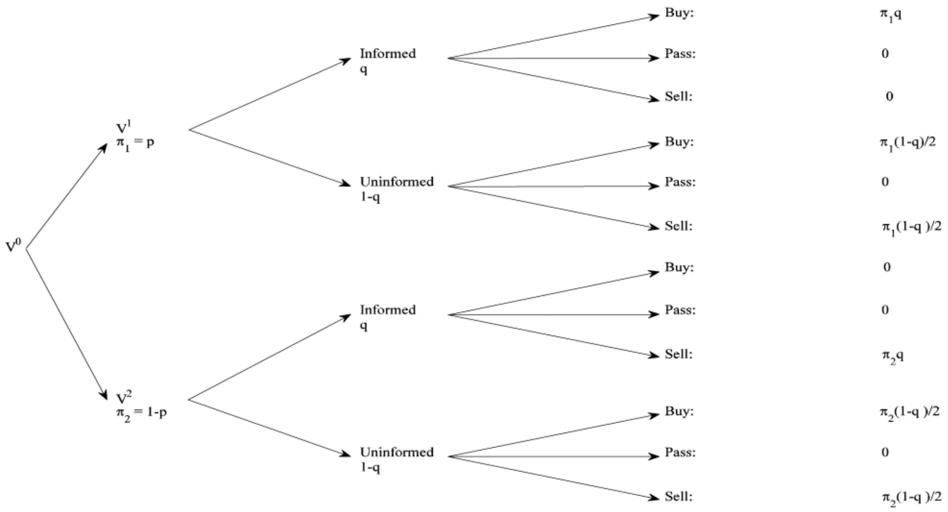

4. Portfolio Strategies

We consider two portfolio strategies: naive diversification and volatility timing. Both strategies are fully-invested, and exclude the possibility of short sales. The naive diversification strategy allocates a weight of

$$w_j = 1/n, \quad j = 1, \dots, n$$

to each portfolio constituent, whereas the volatility timing strategy allocates capital on the basis of a modified version of the minimum variance portfolio.

The problem set-up for the volatility timing strategy is

$$\min_{\mathbf{w} \geq \mathbf{0}} \mathbf{w}'\mathbf{V}\mathbf{w}$$

subject to

$$\sum_{j=1}^{n} w_j = 1,$$

where $\mathbf{w}$ is the vector of portfolio weights, and $\mathbf{V}$ is the variance-covariance matrix of returns. It can be shown that the solution to the problem is

$$w_j = \frac{\psi_j}{\sum_{i=1}^{n} \psi_{j=1}},$$

where ψ_j is the sum of the elements of the jth column of $\mathbf{V}^{-1}$, the inverse of $\mathbf{V}$. However, the volatility timing strategy removes the need to compute the inverse by setting the off-diagonal elements $\mathbf{V}$ of to zero. The elements of the inverse matrix are now simply the reciprocal of the elements of the original matrix, and the solution to the problem is

$$w_j = \frac{1/\sigma_j^2}{\sum_{i=1}^{n} \sigma_i^2}, \tag{16}$$

where σ_j^2 is the variance of the returns to asset j. As with the naive diversification portfolio, the weights of the volatility timing portfolio are non-negative. Both strategies are fully-invested, and both strategies attempt to reduce the high turnover and estimation errors associated with full mean-variance optimization.

We adopt the simplifying assumption that the trades made by portfolio managers do not influence the intra-day dynamics of price.[2] Instead, we assume that orders are good for any size at the opening bid-ask spread, and that portfolio managers place all their orders at the open. This has the additional analytical advantage of separating the influence of private information on transaction costs from its influence (in combination with trading volume) on 'price discovery'. The closing price is used to value positions at the end of each day, which in turn allows calculation of the daily returns to each strategy. Closing prices also provide the information that the volatility timing strategy uses for calculating rolling estimates of daily volatilities—the estimates that in turn determine the desired weights in each asset for the next session. The following algorithms describe the daily activities of the naive diversification and volatility timing strategies.

[2] An interesting enhancement would be to include portfolio managers as part of the trading population, with their orders contributing to daily trading volume—and hence price dynamics. However, the order-splitting strategy of managers needs to be carefully addressed in such a setting.

Algorithm 1: Naive Diversification

Step 1: Revalue account using opening mid-prices

Positions q_j are revalued at the dealer's opening mid-price p_j:

$$\text{account} = \sum_{j=1}^{n} q_j\, p_j + c,$$

where c is the value of short-term cash balances. We allow small temporary negative cash positions, but do not allow strategies to manage leverage strategically.

Step 2: Calculate Desired Positions

The account size is multiplied by 1/*n*, and divided by the dealer's opening mid-price to yield a new desired holding q_j:

$$q_j = \frac{\text{account} \times 1/n}{p_j} \quad j = 1, \dots, n\,.$$

Step 3: Calculate Orders

New orders are calculated as the difference between desired positions q_j and current positions q_j:

$$o_j = q_j - q_j, \quad j = 1, \dots, n.$$

where o_j denotes today's order in asset j.

Step 4: Calculate Expenditure and Income

For buy orders, expenditure is calculated using the dealer's ask price, and for sell orders, income is calculated using the bid price:

$$c_j = \begin{cases} o_j \times \text{Ask}_j & \text{if } o_j > 0 \\ o_j \times \text{Bid}_j & \text{if } o_j < 0 \end{cases}$$

The change in the cash position is the sum of expenditure and income over all assets:

$$\Delta \text{c} = \sum_{j=1}^{n} c_j.$$

Step 5: Revalue Account at Closing Mid-Prices

At the conclusion of intra-day trading, the final dealer quotes are used to calculate closing prices—the final mid-prices for each asset. The account is re-valued, and the daily return to the naive diversification strategy is calculated using

$$r_t = \text{account}_t / \text{account}_{t-1} - 1.$$

Algorithm 2: Volatility Timing

Step 1: Revalue account using opening mid-prices

Positions q_j are revalued at the dealer's opening mid-price p_j:

$$\text{account} = \sum_{j=1}^{n} q_j\, p_j + c,$$

where c is the value of any cash holdings.

Step 2: Calculate Desired Positions

The account size is multiplied by the weights calculated in Equation 16, and divided by the dealer's opening mid-price to yield a new desired holding q_j:

$$q_j = \text{account} \times \frac{1/\sigma_j^2}{\sum_{i=1}^{n} \sigma_i^2} \; / p_j \quad j = 1, \dots, n\,.$$

Step 3: Calculate Orders

New orders are calculated as the difference between desired positions q_j and current positions q_j in each asset:

$$o_j = q_j - q_j, \quad j = 1, \dots, n.$$

where o_j denotes today's order in asset j.

Step 4: Calculate Expenditure and Income

For buy orders, expenditure is calculated using the dealer's ask price, and for sell orders, income is calculated using the bid price:

$$c_j = \begin{cases} o_j \times \text{Ask}_j & \text{if } o_j > 0 \\ o_j \times \text{Bid}_j & \text{if } o_j < 0 \end{cases}$$

The change in the cash position is the sum of expenditure and income over all assets:

$$\Delta \text{c} = \sum_{j=1}^{n} c_j.$$

Step 5: Revalue Account at Closing Mid-Prices

At the conclusion of intra-day trading, the final dealer quotes are used to calculate closing prices—the final mid-prices for each asset. The account is re-valued, and the daily return to the volatility timing strategy is calculated using

$$r_t = \text{account}_t / \text{account}_{t-1} - 1.$$

It now remains to examine large-sample returns and risk-adjusted returns to the two strategies for various market conditions. Each simulation generates 10 years of intra-day trade and closing prices. Each market condition is tested for 1000 simulations of fundamental values.

5. Results

Each cell of Table 2 contains the mean annual return and Sharpe ratio for 1000 multivariate simulations of fundamental values using the single-index model of Section 2. The upper panel contains the results for 2-stock portfolios. The middle panel contains the results for 5-stock portfolios, and the lower panel the results for 10-stock portfolios. Within each panel, the results are split horizontally into those results for the naive diversification strategy, and those for the volatility timing strategy. Vertically, the results are arranged by increasing levels of asymmetric information or probabilities of informed trade. Within each of these sections, an individual cell corresponds (vertically) to the level of trading volume in an illiquid state (10, 50, or 250), and (horizontally) to a level of trading volume in a liquid state (50, 250, 1000). For example, the upper-left cell of the top panel reports a mean return of 10.21%, and a Sharpe ratio of 0.59 for the naive diversification strategy. This corresponds to 1000 underlying simulations of markets in which the probability of informed trade is 0.01, trading volume in the illiquid state is 10, and trading volume in the liquid state is 50. An alternative to the two-state model is a single-state model with constant volume, but we prefer to allow for the possibility of the price discovery mechanism being disrupted at the points where regime shifts occur. The probability of switching is $\zeta = 0.1$, which is intended to reflect our intuition that markets 'remember' the current regime.

For the 2-stock portfolio, an interesting pattern develops as the level of asymmetric information increases from 0.01 through 0.20. For the (10, 50) volume combination, the mean return for the naive diversification strategy is at its highest when q = 0.01. This is to be expected, as the dealer quotes narrow spreads when the probability of adverse selection is low. As the probability of informed trade rises from 0.01 to 0.05, and from 0.05 to 0.10, the naive diversification strategy's mean return falls to 9.85%, and then to 9.31%. The interesting change, however, occurs when the probability of informed trade rises from

0.10 to 0.20: the mean return rises to 9.53%, despite the dealer's wider spreads. A similar pattern occurs for the corresponding Sharpe ratio: 0.59, 0.58, 0.56, and then an increase to 0.60. As this pattern disappears for portfolios with more assets, we suggest that the pattern is linked to the naive diversification strategy's in-built tendency to over-allocate capital to high beta stocks.

The pattern for mean returns is broadly similar for the volatility timing strategy, with apparently little difference between the volatility timing and naive diversification strategies' mean returns under similar market conditions. The most striking difference, however, is in the levels of the Sharpe ratio—the volatility timing strategy consistently produces results approximately 0.20 in excess of those of the naive diversification strategy. In the illiquid volume combination (10, 50), the highest Sharpe ratio of the volatility timing strategy (0.81) occurs when the level of asymmetric information is at its highest. In the most liquid combination (250, 1000), the highest Sharpe ratio occurs when q = 0.05, and falls thereafter. The volatility timing strategy outperforms the naive diversification strategy, not because of its similar mean returns, but because of its lower risk. Efficient price discovery is essential to its success, with the presence of asymmetric information offsetting the impediment to price discovery inherent in low trading volume. That the volatility timing strategy consistently outperforms the naive diversification strategy confirms our view that those strategies that rely on accurate prices, and in turn returns, benefit most from a reasonable level of asymmetric information. In the next section we more formally identify the drivers of portfolio performance.

Table 2. Mean Returns and Sharpe Ratios

2-Stock Portfolio														
		Naive Diversification							**Volatility Timing**					
		50		**250**		**1000**			**50**		**250**		**1000**	
	10	10.21	0.59	9.77	0.58	9.52	0.55		10.17	0.80	9.83	0.79	9.50	0.74
***q*=0.01**	**50**			10.01	0.60	9.96	0.59				10.01	0.80	9.90	0.78
	250					10.03	0.59						10.01	0.80
	10	9.85	0.58	9.70	0.55	9.18	0.58		9.94	0.77	9.69	0.78	9.31	0.79
***q*=0.05**	**50**			9.44	0.58	10.10	0.63				9.53	0.79	10.19	0.85
	250					9.80	0.61						9.88	0.82
	10	9.31	0.56	9.06	0.55	9.31	0.59		9.47	0.76	9.13	0.75	9.58	0.82
***q*=0.10**	**50**			9.34	0.58	9.48	0.60				9.54	0.80	9.42	0.80
	250					10.03	0.60						9.96	0.81
	10	9.53	0.60	9.13	0.56	9.35	0.59		9.59	0.81	9.29	0.76	9.47	0.81
***q*=0.20**	**50**			9.50	0.57	9.38	0.57				9.61	0.77	9.56	0.77
	250					9.74	0.59						9.76	0.80

		5-Stock Portfolio											
		Naïve Diversification						Volatility Timing					
		50		250		1000		50		250		1000	
	10	10.03	0.63	9.90	0.64	9.93	0.64	9.92	1.01	9.99	1.05	9.99	1.06
***q*=0.01**	**50**			10.21	0.64	9.70	0.63			10.16	1.03	9.97	1.05
	250					9.82	0.65					9.86	1.06
	10	9.69	0.63	9.54	0.67	9.21	0.69	9.75	1.04	9.88	1.12	9.62	1.17
***q*=0.05**	**50**			9.46	0.68	9.15	0.69			9.68	1.12	9.55	1.17
	250					9.15	0.68					9.62	1.16
	10	9.17	0.65	9.20	0.69	9.29	0.70	9.45	1.09	9.49	1.18	9.53	1.15
***q*=0.10**	**50**			8.84	0.66	9.20	0.68			9.36	1.10	9.53	1.14
	250					9.22	0.60					9.48	1.01
	10	8.71	0.66	9.08	0.66	8.87	0.65	9.15	1.12	9.48	1.11	9.30	1.10
***q*=0.20**	**50**			9.07	0.61	9.30	0.64			9.54	1.04	9.67	1.07
	250					9.33	0.59					9.54	0.99
		10-Stock Portfolio											
		Naïve Diversification						Volatility Timing					
		50		250		1000		50		250		1000	
	10	9.87	0.69	9.75	0.68	9.89	0.70	9.90	1.21	9.90	1.22	9.98	1.23
***q*=0.01**	**50**			9.77	0.69	9.62	0.69			9.82	1.20	9.77	1.21
	250					9.57	0.70					9.75	1.23
	10	9.59	0.70	9.55	0.76	9.31	0.80	9.85	1.26	9.81	1.35	9.74	1.46
***q*=0.05**	**50**			9.59	0.79	9.38	0.82			9.86	1.40	9.79	1.48
	250					9.36	0.81					9.72	1.44
	10	9.25	0.74	9.15	0.80	9.20	0.80	9.62	1.34	9.59	1.45	9.66	1.46
***q*=0.10**	**50**			9.23	0.79	8.97	0.76			9.67	1.43	9.55	1.42
	250					9.38	0.69					9.72	1.23
	10	8.69	0.76	8.71	0.73	8.57	0.73	9.44	1.42	9.37	1.38	9.38	1.39
***q*=0.20**	**50**			9.13	0.69	8.97	0.68			9.60	1.27	9.56	1.25
	250					9.00	0.62					9.51	1.14

5.1 Data Visualization and Interpretation

Table 2 presents our results in finely-classified samples. While it is clear from the table that the Sharpe ratios of the volatility timing strategy dominate those of the naïve diversification strategy, it is not easy to determine whether strategy type, or some other characteristic of market conditions, is the key driver of portfolio performance. For instance, the probability of informed trade may be important, as may be the simple diversification effect from increasing the number of portfolio constituents. Classification trees, a technique from the nonparametric statistics literature, offer an excellent way of ranking the determinants of portfolio performance, as well as providing a neat visual representation of the data. They are ideally suited to a ranking task, with the data being repeatedly partitioned according to those elements of the sample space that most reduce prediction error. The key reference is Breiman et al. (1984).[3]

The response variable **Y** is predicted using a multivariate set of predictors **X.** In this paper, we consider two response variables—the mean return and the Sharpe ratio of a strategy. The set of predictors includes the strategy type, the number of stocks in the portfolio, the probability of informed trade, and the volume of trade in the illiquid state.

Consider first the two trees for the mean prediction task. Figures 2 and 3 are in fact drawn from one tree, but have been separated to improve legibility. Each sample of observations is represented by an ellipsis or rectangle. The ellipses represent samples that will be divided further into smaller groups; the rectangles, known as the 'leaves' or 'terminal nodes' of the tree, represent the finest partitions of the data. Theoretically, the samples can be partitioned into ever-decreasing samples, until each terminal node contains only one observation, but in practice a tree ceases to be grown (or is 'pruned') according to a statistical or normative criterion. With a view to clarity and parsimony, we terminate the trees using a maximum depth criterion: the number of levels below the initial sample is set to 5, meaning that each of the sub-figures, Figure 2 and Figure 3, has 4 levels. When compared with alternative statistical pruning procedures, we find that this level of detail errs on the side of parsimony—further nodes, by definition, improve the in-sample predictive accuracy of the tree, but do so with increased risk of over-fitting.

[3] Software-based tutorials include Martinez and Martinez (2008) and Torgo (2011).

There are 144,000 observations in the full sample, which correspond to the 144 cells of Table 2. The pooled mean return is 9.55% for these observations. The procedure examines all possible partitions across the predictor variables, and chooses the best binary split— defined as the partition that most reduces the total mean-squared-error of the tree.

Formally, the mean of the full sample is defined by

$$y = \frac{1}{n} \; y_i \, ,$$

and the mean-squared-error by

$$R = \quad y_i - y \;\; ^2 .$$

After the first split, there are two nodes, each with its own y and R. Denote these within-node squared errors by $R_1$and R_2. After the first split, the mean-squared-error of the tree is the sum of $R_1$and R_2.

The first partition of the mean returns classification tree splits the full sample into two groups. The first contains observations in which the probability of informed trade $q \in \{0.01, 0.05\}$, and the second contains observations in which $q \in \{0.10, 0.20\}$. These are the nodes at the top of the trees of Figures 2 and 3. Interpreting these two new samples as 'low' and 'high' asymmetric information samples, the low asymmetric information tree of Figure 2 has a sample mean of 9.76%, with 72,000 observations. The high asymmetric information tree also contains 72,000 observation, but with a mean return of 9.35%. The 0.41% fall in mean returns represents a flow of wealth from uninformed to informed traders.

We next consider the low asymmetric information and high asymmetric information trees. For each tree, we describe a particular path down the tree. The highest mean return of Figure 2 is 9.97%, represented by the second terminal node from the left. The first partition of the tree divides observations into probabilities of informed trade of q = 0.01 and q = 0.05. In markets with the lowest probability of informed trade (q = 0.01), mean returns are 9.89%, 0.27% higher than the mean returns generated by markets with q = 0.05. Continuing down the q = 0.01 branch of the tree, the next most important driver of performance is the number of portfolio constituents, with those portfolios containing 2 or 5 stocks generating higher returns than those with 10 stocks. An explanation for this phenomenon could be that the naive diversification strategy allocates equal levels of capital to high-volatility and low-volatility assets, thus boosting mean returns when the number of portfolio constituents is small. More clearly, in the next partition, mean returns are higher when trading volume is high in the illiquid state. Note that this path down the tree does not distinguish the returns to the naive diversification strategy from those to the volatility timing strategy. The story is different, however, in the high asymmetric information tree.

Figure 3 demonstrates that, conditional on the probability of informed trade being 10% or 20%, the next partition that most reduces the tree's mean-squared-error is strategy type. The returns to the naive diversification strategy are 9.18%, whereas they are 9.52% for the volatility timing strategy. This would suggest that the naive diversification strategy generates a higher turnover of trade than does the volatility timing strategy—a feature that most impacts on performance when bid-ask spreads are wide. Continuing down the volatility timing strategy path, the next partition is according to the volume in the illiquid state, with volumes of 50 or 250 generating higher returns than a volume of 10. The final partition, branching to the farthest-right terminal node (9.66% across 6,000 observations), further distinguishes illiquid state trading volume of 50 from trading volume of 250. In sum, in markets characterized by high levels of asymmetric information, a volatility timing strategy applied under conditions of high overall liquidity generates on average mean returns of 9.66% p.a.

We conclude this section by examining the classification tree in which the response variable is the Sharpe ratio of the volatility timing strategy (Figure 5). We describe the path that leads to the highest Sharpe ratios. The root node of Figure 5 contains 72,000 observations, with a mean Sharpe ratio of 1.07. The next partition is with respect to the number of portfolio constituents, with the Sharpe ratios of 5 or 10 stock portfolios, substantially exceeding those of 2-stock portfolios (1.21 versus 0.79). Of those portfolios with 5 or 10 stocks, the next partition distinguishes the 5-stock portfolios from the 10-stock portfolios. On average, 10-stock volatility timing portfolios generate Sharpe ratios of 1.33, 0.24 higher than 5-stock portfolios. It is interesting to note, however, that the terminal nodes across the tree are partitioned according to the probability of informed trade. For 2-stock portfolios there appears to be little difference between the average Sharpe ratios in each sample, with terminal nodes containing average Sharpe ratios of 0.78, 0.79, 0.81, and 0.78. However, for better-diversified portfolios, the optimal partition splits the observations by probabilities of informed trade of 1%, and probabilities of informed trade of greater than 1%. Even though higher levels of private information lead to wider spreads and higher transaction costs, they lead to higher Sharpe ratios for both the 5-stock and 10-stock portfolios. This 'price discovery' effect is most pronounced in the 10-stock portfolios, with an improvement in the Sharpe ratio from 1.21 to 1.37 when the probability of informed trade exceeds 1%.

6. Conclusions

We develop a framework in which multi-asset fundamentals are mapped into binomial processes compatible with the Glosten and Milgrom (1985) sequential trade model. Intra-day price dynamics are generated by dealers' Bayesian updating equations, with closing prices determined by the average of the final bid and ask prices of each session. The degree to which closing prices track fundamental value is determined by the joint interaction between private information and trading volume. Higher levels of private information reduce mean returns, as dealers widen spreads to compensate for the losses incurred from informed trade. But private information also helps to improve the price discovery process, thus improving the risk-adjusted returns of strategies that rely on accurate volatility estimates.

We use nonparametric classification trees to identify and rank the determinants of portfolio performance. Mean returns are primarily driven by the probability of informed trade, whereas the strategy type - naive diversification or volatility timing - is the key driver of risk-adjusted returns. This suggests that the higher Sharpe ratios of the volatility timing strategy arise because of its objective of minimizing risk; this does not appear to sacrifice mean returns. The diversification effect from increasing the number of portfolio constituents is the next most important driver of risk-adjusted returns, with the highest Sharpe ratios of both strategies occurring in the 10-stock portfolios. We note the interesting dominance of the volatility timing Sharpe ratios in markets when the probability of informed trade is greater than 1%. Indeed, looking down the columns of Table 2, it is evident that the lower mean returns associated with wider spreads are often accompanied by higher Sharpe ratios, there being an apparent 'optimal' level of asymmetric information, beyond which Sharpe ratios decline. These declines occur as higher levels of transaction costs begin to dominate improvements in the price discovery mechanism.

With regard to extensions and future research, we have deliberately designed the framework with flexibility in mind. We have used a single-index model, but envisage more elaborate factor models in the data generation stage. The recombining tree structure of our sequential trade model allows for stochastic news arrivals, whilst keeping the dealer's updating task manageable. We would maintain the informational advantage of the insiders during the 'no news' days, thus making a distinct contribution to the literature.

References

1. Breiman, L., Friedman, J., Olshen, R., Stone, C., 1984. Classification and regression trees. Wadsworth, Inc., New York.
2. Brunnermeier, M. K., 2001. Asset pricing under asymmetric information. Oxford University Press, Oxford.
3. Cox, J. C., Ross, S. A., Rubinstein, M., 1979. Option pricing: a simplified approach.
4. Journal of Financial Economics 7, 229–263.
5. DeMiguel, V., Garlappi, L., Uppal, R., 2009. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? Review of Financial Studies 22, 1915–1953.
6. Easley, D., O'Hara, M., 1987. Price, trade size, and information in securities markets. Journal of Financial Economics 19, 69–90.
7. Easley, D., O'Hara, M., 1992. Time and the process of security price adjustment. Journal of Finance 47, 577–605.
8. Einrich, L. J., Piedmonte, M. R., 1991. A method for generating high-dimensional multivariate binary variates. The American Statistician 45, 302–304.
9. Glosten, L., Milgrom, P., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71–100.
10. Ho, T., Stoll, H. R., 1981. Optimal dealer pricing under transactions and return uncertainty. Journal of Financial Economics 9, 47–73.
11. Kirby, C., Ostdiek, B., 2012. It's all in the timing: Simple active portfolio strategies that outperform nave diversification. Journal of Financial and Quantitative Analysis 47 (02), 437–467.
12. Kyle, A., 1985. Continuous auctions and insider trading. Econometrica 53, 1315–1335.
13. Martinez, W. L., Martinez, A. R., 2008. Computational Statistics Handbook with MATLAB. Chapman & Hall/CRC, Boca Raton,FL.
14. O'Hara, M., 1995. Market Microstructure Theory. Blackwell, Cambridge, MA.
15. Rebonato, R., Jackel, P., 1999. The most general methodology to create a valid correlation matrix for risk management and option pricing purposes. Quantitative Research Centre of the NatWest Group, 1–12.
16. Torgo, L., 2011. Data mining with R: Learning with case studies. Chapman & Hall/CRC, Boca Raton, FL.

Figure 2. Mean Returns: q ∈ {0.01, 0.05}

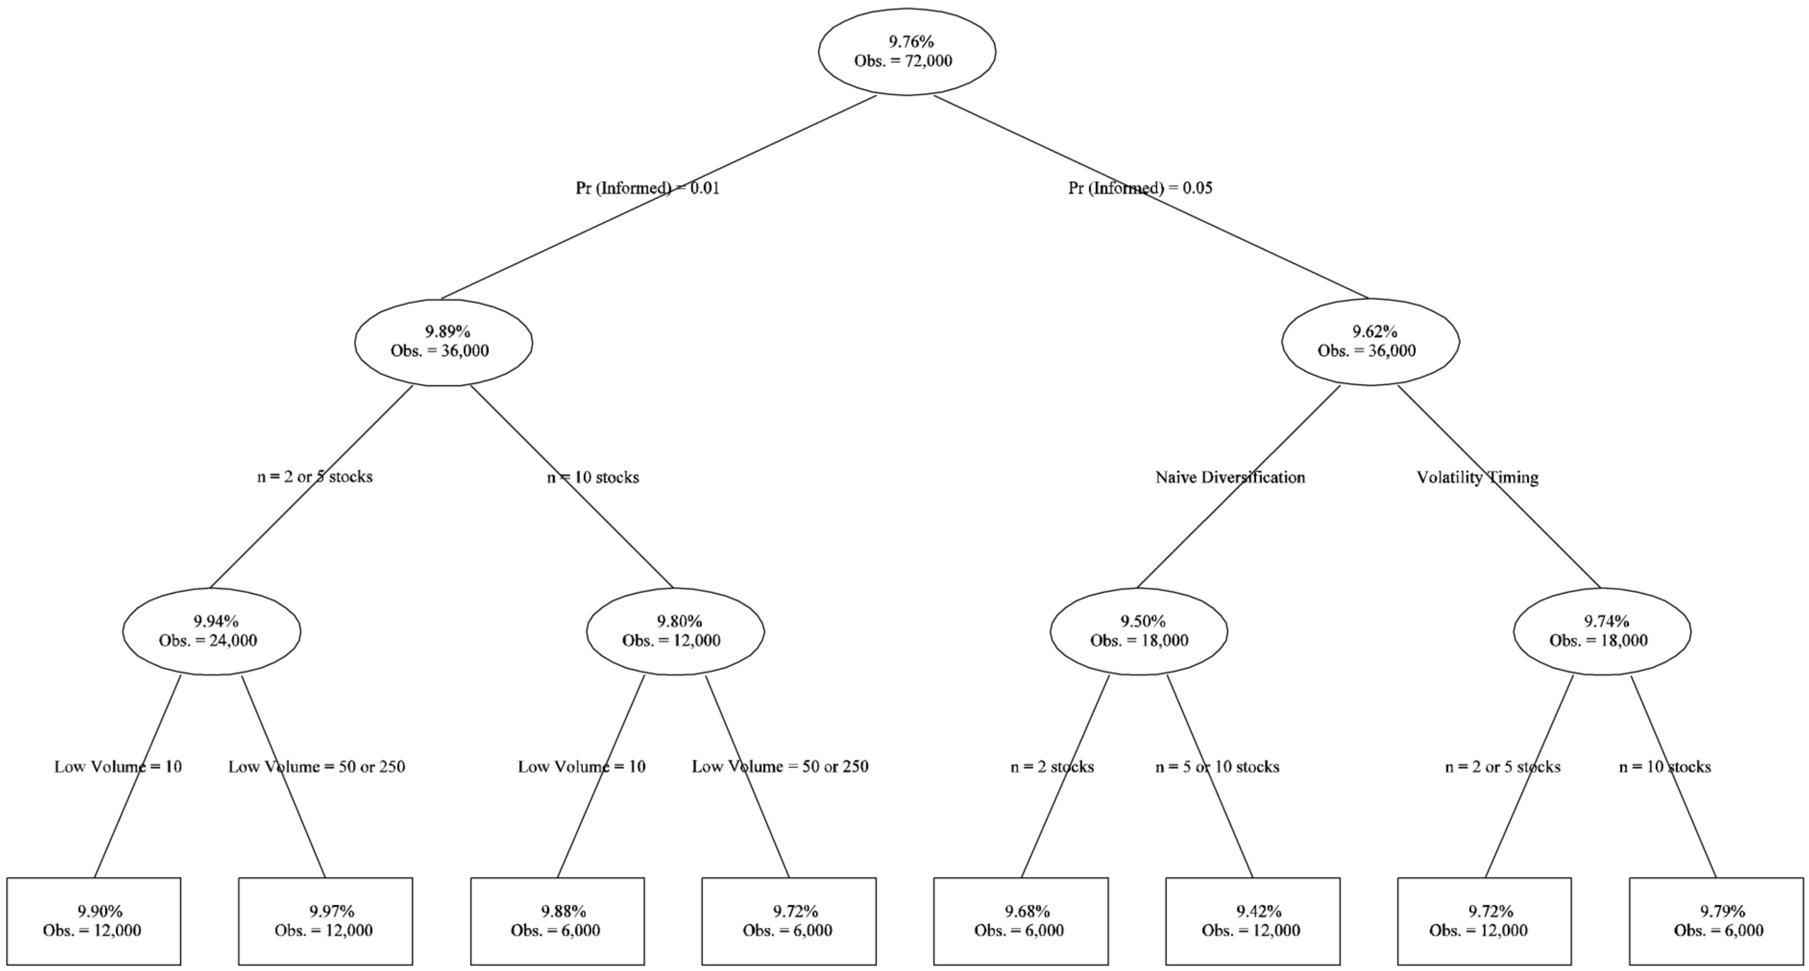

Figure 3. Mean Returns: q ∈ {0.10, 0.20}

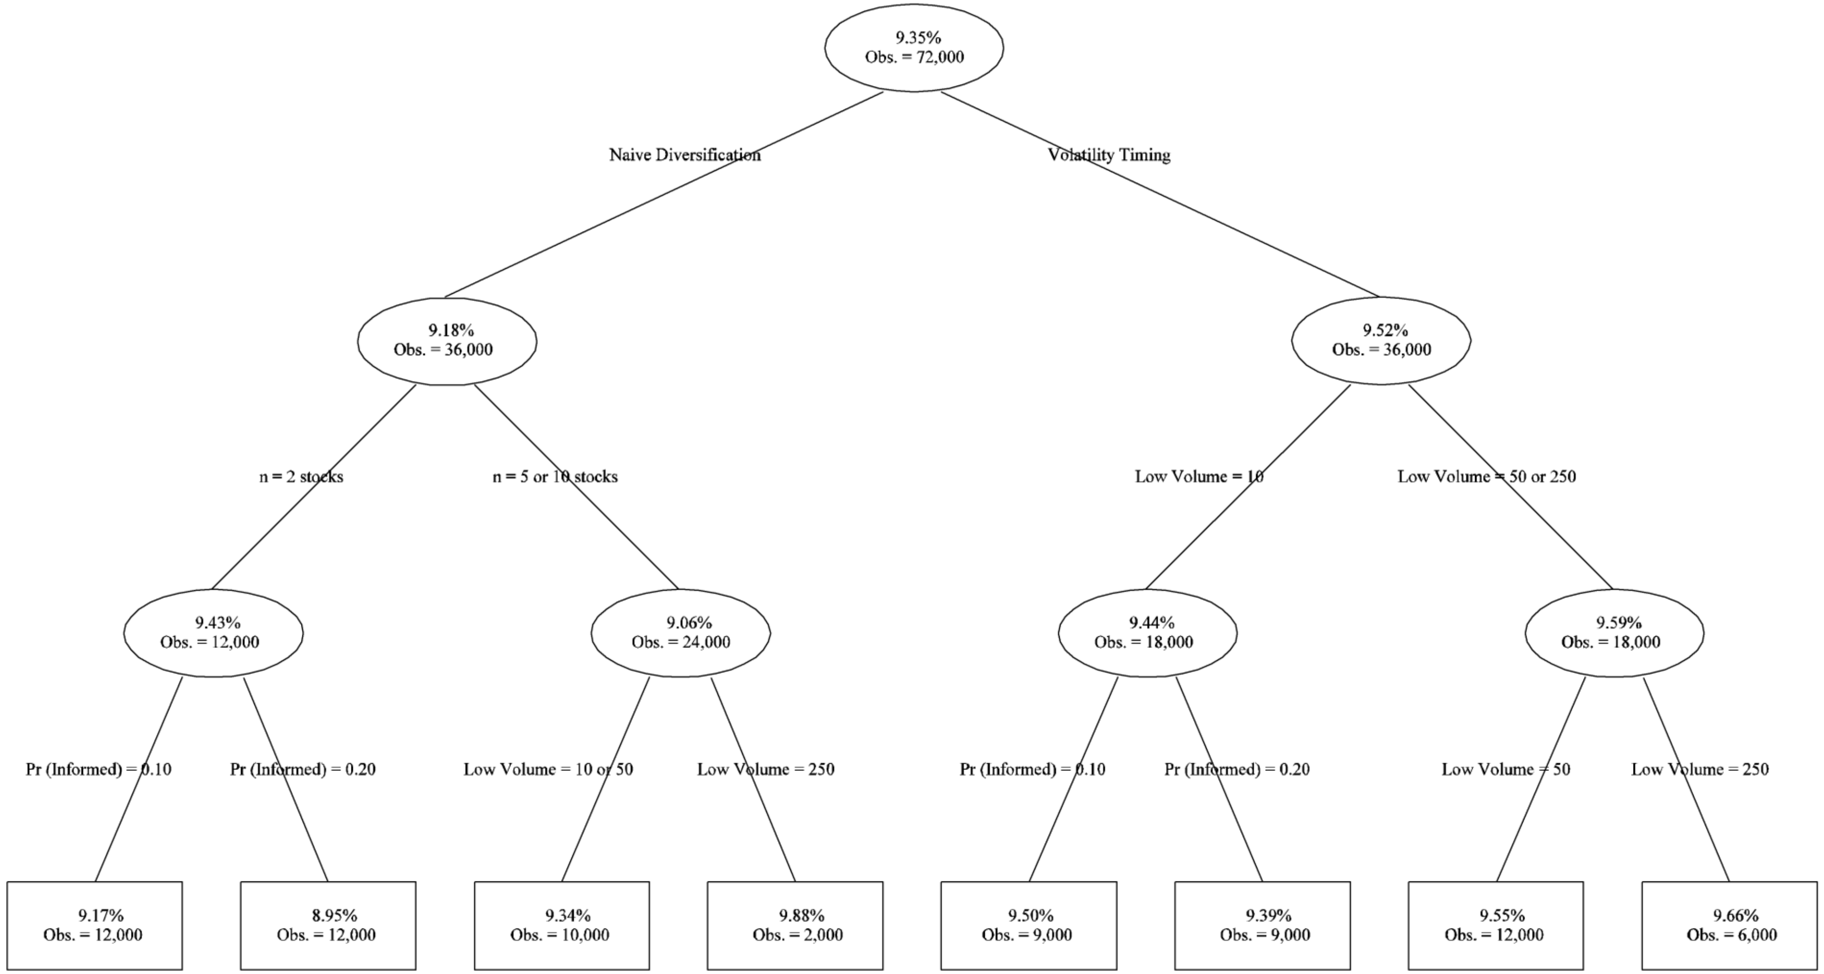

Figure 4. Sharpe Ratios: Naïve Diversification

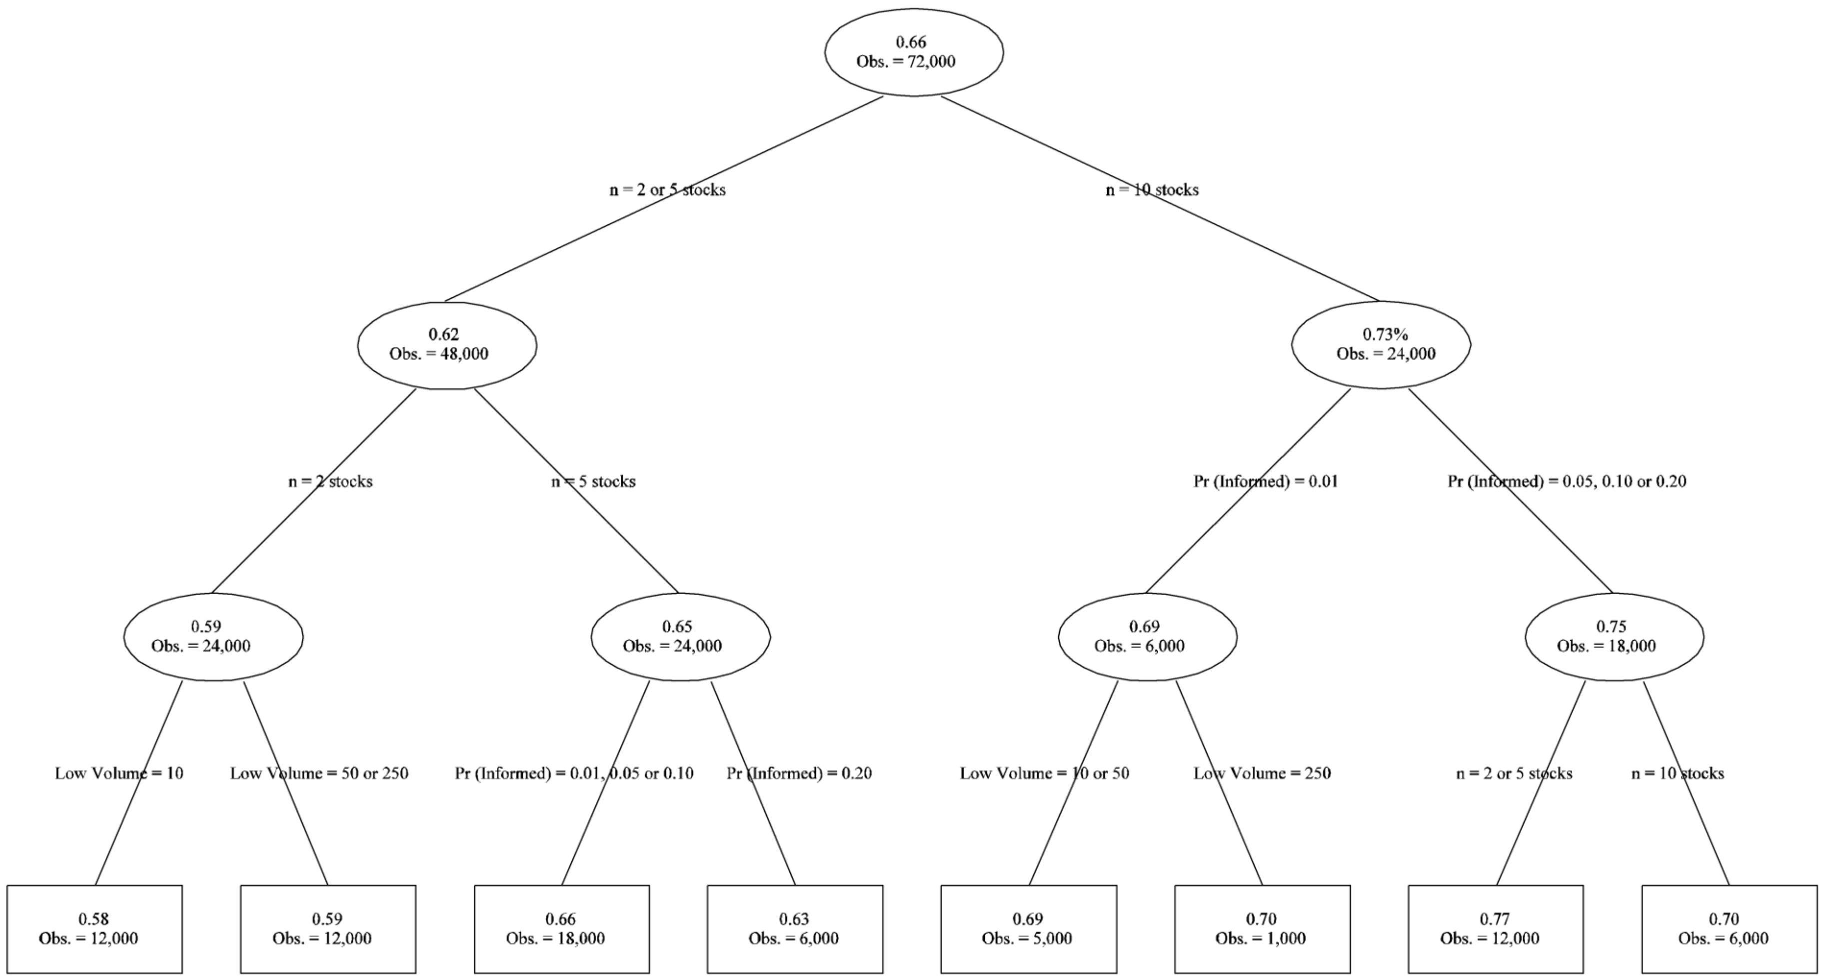

Figure 5. Sharpe Ratios: Volatility Timing