Volume 5 • Number 1
Fall 2015

JURF
Journal of Undergraduate
Research in Finance

Not Another Market Timing Scheme! : Detecting Type I Errors with "Good Deal" Bounds

Michael Lachanski[1]

[1]The idea for the paper comes from Steven Pav, who supplied the analysis of market timing schemes and directed me to Quant Finance Stack Exchange post explaining the good deal bound. This project was inspired, supervised, and funded by Professor Jianqing Fan and Princeton University's Program in Applied and Computational Mathematics. I am grateful to Professors Frank Fabozzi, Erhan Cinlar, Jonathan Godbey and an anonymous referee for their helpful comments.

Abstract

Unlike traditional financial econometrics applications, the machine learning for stock market prediction literature (especially that which relies on textual analysis) typically evaluates algorithms by their ability to forecast the direction, but not the magnitude of asset price movements. Combining the market timing model of Pav (2012) and the "good deal" bounds of Ross (2005), we give economic content to these results by obtaining explicit good deal bounds on market direction predictability. In a numerical example, we show that the market predictability results reported in Bollen et al. (2011) and Vu et al. (2012) violate the good deal bound and are likely to be overstated.

I Introduction: Motivation, Definitions, and Baseline Results

Nassirtoussi et al. (2014) summarize many studies suggesting one can use text mining, or the statistical analysis of natural language data to predict asset price movements with high accuracy. This essay reviews the results of the statistical natural language processing for market prediction literature, reviews the general representation of market timing ability first given by Pav (2012), reviews bounds on asset performance constructed by Ross (2005) and combines these latter two results to bound market timing ability. The key is that the result presented in Pav (2012) allows one to map market timing ability to Sharpe Ratio (SR hereafter)[1] and Ross (2005) bounds SR in terms of observable parameters.

I.1 Some Definitions

For a given portfolio h (where we make h precise below), the SR of that portfolio in the next period is defined as:

$$SR\left(h\right) \coloneqq \frac{\mathbb{E}_t^2\left[r_{t+1}^h\right] - r_f}{\sigma_t\left(r_{t+1}^h - r_f\right)}$$

where $\mathbb{E}\left[r_{t+1}^h\right]$ is the expected return of portfolio h in the next period given all the information up to time t, r_f is the risk-free rate and $\sigma\left(r_{t+1}^h - r_f\right)$ is the expected standard deviation of the excess return portfolio h given all information up to time t. We will use this notation throughout the essay but drop the t subscripts from our expectation operators. The SR was developed by Sharpe (1966) as an ad hoc measure to assess the performance of asset managers and is perhaps the most common method of doing so today. As one might surmise from the discussion, the SR can be an economically meaningful quantity, but we can see that "upside" risk is penalized in the denominator as much as "downside" risk, which means that higher SRs are not unambiguously better than low SR; in particular, Brunnermeier (2015) provides an example in a portfolio with low SR has higher return in every state of the world than a portfolio with high SR. We will impose additional restrictions on our market timing model so that portfolios with higher SRs are unambiguously better choices than portfolios with lower SRs. In particular, the bound Ross (2005) constructs is based on an extension of the No Arbitrage (NA hereafter) bound on portfolio performance called the no "good deal" (GD hereafter) bound. The "good deal" bound uses a minimal characterization of the representative investor to further restrict the statistical performance of a portfolio.

Taking purported accuracy of predicting markets as a statistical parameter, we can call it a Type I error when the true ability to predict the market is lower than the reported accuracy and a Type II error when the true ability to predict the market is higher than the reported accuracy. If reported accuracy is higher

[1] In response to numerous requests, we have added a nomenclature section to this essay.

than our bound it is either a Type I error or evidence against the assumptions giving rise to the GD bound.[2]

I.2 Some Stylized Facts

Nassirtoussi et al. (2014) reviews several dozen studies on text mining for stock market prediction. In their essay, they also provide the maximal up-down accuracy of several studies where accuracy is defined as in (1) below:

$$Accuracy := \frac{\#\text{ Correct Predictions}}{T-1} \tag{1}$$

where T is the total number of days in the sample and the number of correct predictions refers to the number of times the model was able to correctly anticipate whether day t's asset price was greater or lower than day $t-1$'s asset price with $t = 2, \ldots, T$. While statistical models in financial econometrics are rarely evaluated on the basis of their ability to predict the up-down sequence of the stock market indices or equity prices, this benchmark is fairly common in the textual analysis and machine learning for stock market prediction literature. This is not without justification because much of the machine learning work in text analytics for stock market prediction is meant to be deployed in high-frequency trading strategies for which the primary criterion of success is the ability to predict the up-down tick sequence over a short time frame. We present the maximal predictive accuracies of several market timing strategies assessed in Nassirtoussi et al. (2014) in Table 1.

While the wide variety of financial instruments and methodologies used in the studies surveyed by Nassirtoussi et al. (2014) prevent us from conducting a formal meta-analysis, the rank ordering of market predictability by study highlights Bollen et al. (2011) as the candidate most likely to be a statistical outlier.

II Converting Timing Ability to Sharpe Ratio

In this section, we review the market timing model presented by Pav (2012). In this model, we assume that our trader has unit leverage and is in the market at all times.

II.1 General Market Timing Model

Let r_m give the daily returns of portfolio h.[3] Then let $r_m - r_f = r_t$ is the excess return of the portfolio. We will assume throughout this essay that the risk-free rate is approximately 0 for convenience. Our ability

[2] All code for this essay can be found at `https://github.com/mlachans/NAMTS`.

[3] While we can choose the unit of time arbitrarily, all of the up-down predictability studies shown in Table 1 use the day as their frequency. SR typically uses annual units, and so we will need to convert out market timing ability to an annualized SR.

Reference	Maximal Out-of-Sample Predictability	Financial instrument(s)
Wuthrich et al. (1998)	53%	Hsi Index
Soni et al. (2007)	56.2%	Dutch oil and natural gas stocks
Schumaker et al. (2012)	58.2%	S&P 500 stocks
Mahajan et al. (2008)	60%	Sensex (cap weighted index of 30 stock)
Zhai et al. (2007)	70.1%	BHP (a large Australian mining company)
Hagenau et al. (2013)	65.1%	Select stocks from the UK and Germany
Rachlin et al. (2007)	82.4%	CSCO, EBAY, MSFT, TEVA, YHOO[a]
Vu et al. (2012)	84.62%	AAPL, GOOG, MSFT, AMZN
Huang et al. (2010)	85.32%	Foreign exchange markets
Bollen et al. (2011)	86.7%	DJIA

Table 1: Select results compiled from Table 5 in Nassirtoussi et al. (2014) and Bollen et al. (2011) with up-down sequence predictions on financial instruments at the daily or higher frequency using textual machine learning are presented. Words in all capital letters refer to tickers.

[a] Accuracy decreases when text is included.

to time the market comes from a model which we will say provides us a signal s_{t-1}. In this case:

$$s_{t-1} = \begin{cases} \mathrm{sgn}\,(r_t) & \text{with } p \\ -\mathrm{sgn}\,(r_t) & \text{with } 1-p \end{cases} \tag{2}$$

Notice that, in this model, we can predict the direction but not the magnitude of the portfolio's movements. If the probability of our model making a correct prediction p is far below 0.5 we can simply bet against it (i.e. $p = 0$ would imply we could always negate the signal to have a perfect predictor of the stock's up-down pattern).[4] Thus, our analysis will be cleaner if we define a parameter representing our statistical edge $g \in \left[0, \frac{1}{2}\right]$ such that:

$$p - \frac{1}{2} = g \tag{3}$$

where, to be concrete, we can take Mahajan et al. (2008) as claiming that their algorithm applied to the excess return of the Sensex has $g = \frac{1}{10}$. The portfolio in our model is a binomial tree with probability q of an increase and probability $1-q$ of a decrease. We show the market timing strategy in Figure I.

In this model, we have $\mathbb{E}\left[r_t | r_t > 0\right]$ probability p, $\mathbb{E}\left[-r_t | r_t > 0\right]$ with $(1-p)$, $\mathbb{E}\left[-r_t | r_t < 0\right]$, with probability p and $\mathbb{E}\left[r_t | r_t < 0\right]$ with probability $(1-p)$. We can put these four cases together to see that our expected return is:

$$\begin{aligned} \mathbb{E}\left[r_t | r_t > 0\right] p + \mathbb{E}\left[-r_t | r_t < 0\right] p + \mathbb{E}\left[r_t | r_t < 0\right](1-p) + \mathbb{E}\left[-r_t | r_t > 0\right](1-p) &= p\mathbb{E}\left|r_t\right| - (1-p)\,\mathbb{E}\left[|r_t|\right] \\ &= (2p-1)\,\mathbb{E}\left[|r_t|\right] \\ &= 2g\mathbb{E}\left[|r_t|\right] \end{aligned}$$

where the right hand side of the above is our expected return for the numerator of the SR.

II.2 The SR of a Market Timer

Now, we must obtain the denominator of the SR. Recalling that $\mathbb{V}\left[X\right] = \mathbb{E}\left[X^2\right] - \mathbb{E}\left[X\right]^2$, we can use the assumption that we are always in the market to note that our expected squared return sequence must simply be the expected squared return sequence of the market $\mathbb{E}\left[r_t^2\right]$. Pav (2012) concludes that the SR can be written in closed form as follows:

[4]In practice, of course, this means that studies exhibiting predictability significantly less than 50% are rarely published.

A Simple Market Timing Strategy

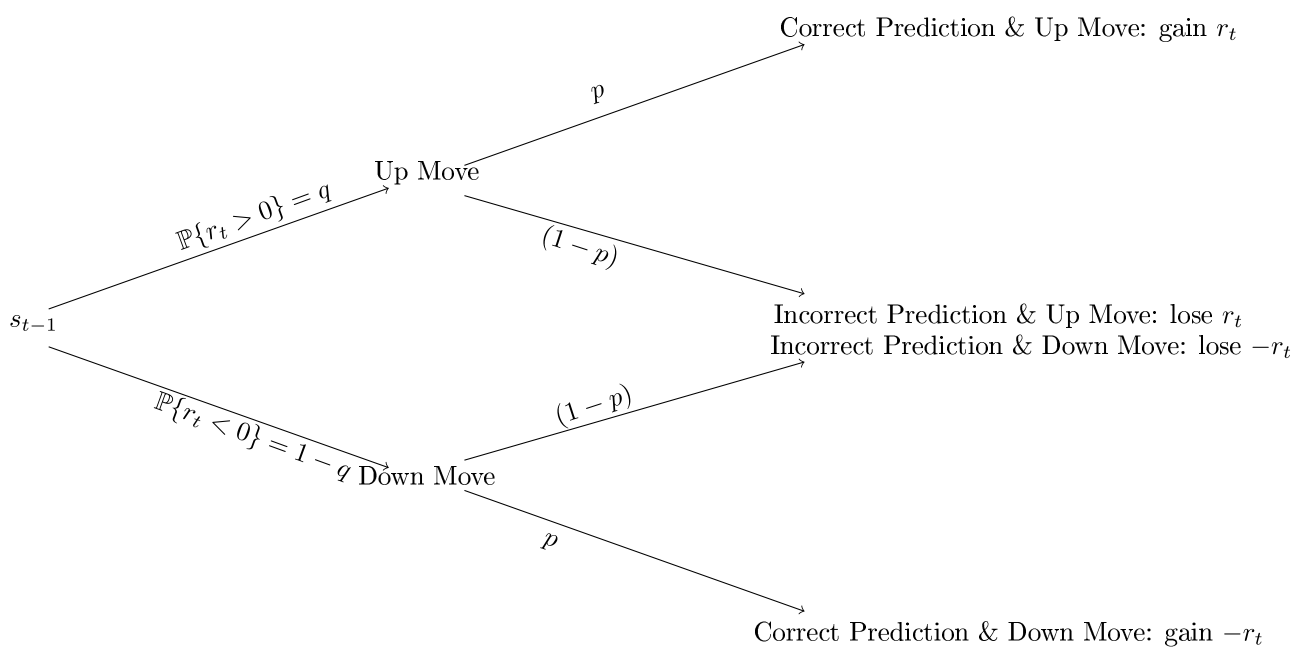

Figure I: In our market timing model the market is a binomial tree with probability q of an increase and $1 - q$ of a decrease. Our market timing agent is always in the market and our probability of a correct forecast of the up-down pattern is $p = \frac{1}{2} + g$. Our probability of an incorrect guess is $1 - p = \frac{1}{2} - g$

$$SR_{\text{Market Timer}} = \frac{2g\mathbb{E}\left[|r_t|\right]}{\sqrt{\mathbb{E}\left[r_t^2\right] - 4g^2\mathbb{E}\left[|r_t|\right]^2}} = \frac{g}{\sqrt{\frac{\mathbb{E}\left[r_t^2\right]}{4\mathbb{E}\left[|r_t|\right]^2} - g^2}} = \frac{g}{\sqrt{\kappa - g^2}} \tag{4}$$

where in the last equality he defines $\kappa := \frac{\mathbb{E}\left[r_t^2\right]}{4\mathbb{E}\left[|r_t|\right]^2}$. We can estimate κ from our portfolio's daily excess returns:

$$\frac{\frac{1}{T}\left[\sum_{t=1}^T r_t^2\right]}{4\left[\frac{1}{T}\sum_{t=1}^T |r_t|\right]^2} = \frac{T\left[\sum_{t=1}^T r_t^2\right]}{4\left[\sum_{t=1}^T |r_t|\right]^2} = \hat{\kappa}_{\text{emp}}$$

so that when r_t is a standard normal, we obtain exact estimates of SR. Specifically, we can use the fact that, for a standard Wiener process W_t:

$$\mathbb{E}\left[|W_t|\right] = \int_0^\infty \frac{2x}{\sqrt{2\pi t}} e^{-\frac{1}{2}x^2} dx = \sqrt{\frac{2t}{\pi}} \text{ and } \mathbb{E}\left[W_t^2\right] = \mathbb{V}\left[W_t\right] = t$$

so that, letting the variance of one day's equity return movement be given by t we have:

$$\kappa\left(t\right) = \frac{t}{4\left(\frac{2t}{\pi}\right)} = \frac{\pi}{8} \approx 0.39$$

where the form above shows that κ has no obvious interpretation. It does not necessarily depend on the variance as one might assume from the numerator of its definition; in the case of a Gaussian distribution, κ is independent of the variance. It appears that κ is increasing in the kurtosis of the distribution it is drawn from. To see this, we can simulate large samples from t-distributions with different degrees of freedom κ is increasing the heaviness of the tails of the t-distribution. In Table 2, we use Monte Carlo simulation to estimate κ for four samples of 10,000 with our degrees of freedom parameter taking on values 4, 4.5, 5 and 10.

III Neoclassical Finance: Bounding Market Predictability

In this section, we will present the Hansen-Jagannathan bound, published results on relative risk aversion and show how these suggest an upper bound on asset performance defined in terms of SR. Our presentation of the Hansen-Jagannathan bound derivation is adapted from Brunnermeier (2015). In equilibrium, the pricing kernel is a function of the marginal utility of risky assets which can, in principle, be estimated from either experiments or aggregate data. The easiest way to bound the pricing kernel is to assume that the marginal investor has a utility function parameterized by relative risk aversion. Then, we can use bounds on relative risk aversion to bound the volatility of the pricing kernel, which in turn, bounds maximal achieveable SR. Assets that deliver SRs greater than than the volatility of the pricing kernel over any significant length

Degrees of Freedom	$\hat{\kappa}$	Asset	$\hat{\kappa}$
4	0.49	AAPL	0.48
4.5	0.48	GOOG	0.54
5	0.47	MSFT	0.47
10	0.41	AMZN	0.49
∞(Gaussian)	0.39	S&P 500	0.52

Table 2: This table contains the results of our Monte Carlo simulation for the t-distribution under different degrees of freedom as well as empirical estimates of $\hat{\kappa}$ for each company from 2010 to 2014. It suggests that κ is increasing in the heaviness of the tails of the distribution. All data is obtained from Google Finance.

of time are called "good deals" and are generally rare. We will be working in the Arrow-Debreu general equilibrium one-period model in which each asset's returns are identically distributed each period and assets returns are independent in each period, but assets returns within a period are not independent of each other. We can make a "Debreu" style multi-period model by simply treating each unit of calendar time (day/week/month/year) as an instance of the one-period model and assuming that the security structure is fixed across time. This may seem unrealistic but the studies cited in Table 1 actually fixed their portfolios through time, in most cases focusing on predicting the up-down sequence of prices for a single asset. In general, a major weakness of this work is that even a single asset might exhibit time dependence so that each period's return is not independent across time. The general "Arrow" style multi-period model can handle all instances of time dependence, but as one might suspect it is very complicated. It is especially difficult to connect the full multi-period discrete time model with data and so we leave this extension as a problem for future work.

III.1 The Arrow-Debreu Model: Asset Pricing in Discrete Time

We assume there are S states of the world $s = 1,\ldots,s = S$ and that we have J assets $j = 1,\ldots,J$. Each asset j pays x^j_s in state of the world s so that:

$$x^j = \begin{bmatrix} x^j_1 \\ x^j_2 \\ \vdots \\ x^j_{S-1} \\ x^j_S \end{bmatrix}$$

where each state s has probability $\pi_s > 0$ of occurring such that $\sum_{s=1}^{S} \pi_s = 1$. All assets are linearly independent of each other by assumption. We can put all of these assets together to define a "security structure" with each row defining a state and each column defines an asset:

$$X = \begin{bmatrix} x^1_1 & x^2_1 & \cdots & x^{J-1}_1 & x^J_1 \\ x^1_2 & x^2_2 & \cdots & x^{J-1}_2 & x^J_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ x^1_{S-1} & x^2_{S-1} & \cdots & x^{J-1}_{S-1} & x^J_{S-1} \\ x^1_S & x^2_{S-1} & \cdots & x^{J-1}_S & x^J_S \end{bmatrix}$$

Markets are complete if X has full rank (e.g. $J = S$) and we will assume market completeness throughout.

Ross (2005) suggests that extending our model to include incomplete markets is unlikely to change the results of this exercise. Now, we can define a portfolio $h \in \mathbb{R}^J$ as simply a fixed quantity of assets which specifies a payoff hX in all states of the world.

III.1.1 Asset Pricing Formulas: a Review

Brunnermeier (2015) show that under NA there must be a linear combination of assets for each state, i.e. a portfolio, that pays $1 for that state and 0 for all other states and the price of this portfolio is positive. We call the price of the portfolio corresponding to this payoff the state price q_s. Given NA, we can obtain the price of any asset as simply the cash flows of that asset in each state multiplied by the state price for one dollar as in (5) below:

$$p_j = \sum_{s=1}^{S} q_s x_s^j \tag{5}$$

where p_j is the market price of asset x^j. From (5), we can derive a number of asset pricing formulas in discrete time. The most common asset pricing formula represents the price of any asset as $p = \mathbb{E}[mx]$ where m is called a stochastic discount factor. We obtain this below:

$$p_j = \sum_{s=1}^{S} \pi_s \left(\frac{q_s}{\pi_s}\right) x_s^j \equiv \sum_{s=1}^{S} \pi_s m_s x_s^j = E\left[x^j \cdot m\right] \tag{6}$$

where $m_s \equiv \frac{q_s}{\pi_s}$. Our stochastic discount factor (SDF hereafter) is:

$$m \equiv \begin{bmatrix} m_1 \\ m_2 \\ \vdots \\ m_{S-1} \\ m_S \end{bmatrix}$$

We can rewrite the right hand side of (6) by using the fact that for a risk-free bond we must have, by definition, $\forall s, \quad x_s^b = 1$. Thus, using (6) :

$$\begin{aligned} p_b = \mathbb{E}[m] = \frac{1}{R^f} \quad &\text{and} \quad p_j = \mathbb{E}\left[mx^j\right] = \mathbb{E}\left[x^j\right]\mathbb{E}[m] + \sigma\left(m, x^j\right) \\ \Rightarrow p_j \quad &= \quad \frac{\mathbb{E}\left[x^j\right]}{R^f} + \sigma\left(m, x^j\right) \end{aligned}$$

Now, note that if we work with gross returns $R^j \equiv \frac{x^j}{p_j}$ then we can always write: $\mathbb{E}\left[mR^j\right] = 1$. Since

$R^f = \frac{1}{\mathbb{E}[m]}$, we have:

$$\begin{aligned} \mathbb{E}\left[m \cdot \left(R^j - R^f\right)\right] &= 0 \\ \Rightarrow 0 &= \mathbb{E}\left[m\right]\mathbb{E}\left[R^j - R^f\right] + \sigma\left(m, R^j\right) \\ \Rightarrow \mathbb{E}\left[R^j\right] - R^f &= -\frac{\sigma\left(m, R^j\right)}{\mathbb{E}\left[m\right]} \end{aligned} \tag{7}$$

Thus, in the one period NA model, all portfolio returns come from the portfolio having negative covariance with the stochastic discount factor (the risk premium) and the time value of money. By imposing additional restrictions on (7) we can obtain the familiar Sharpe-Lintner-Treynor Capital Asset Pricing Model; as one might expect, we can also extend (7) to the multi-period case to obtain the Ross factor pricing model or Consumption Capital Asset Pricing Model.

III.1.2 A Potentially Helpful Detour: Defining Risk-Neutral Probabilities in Discrete Time

We can use (6) to see that the price of a the risk -free bond must be $p_b = \sum_{s=1}^{S} q_s = \frac{1}{R^f}$. With this fact, we can also derive the equivalent martingale measure in discrete time, which relates quantities (specifically risk-neutral probabilities) in the continuous time asset pricing framework with the one-period discrete time model. We can rewrite (5) as:

$$p_j = \sum_{s=1}^{S} q_s x_s^j = \frac{1}{R^f}\sum_{s=1}^{S} R^f q_s x_s^j = \frac{1}{R^f}\sum_{s=1}^{S}\left(\frac{q_s}{\sum_{s=1}^{S} q_s}\right) x_s^j$$

Now, since NA implies that $q_s > 0$ and we know that $\sum_{s=1}^{S}\left(\frac{q_s}{\sum_{s=1}^{S} q_s}\right) = 1$ it must be that $\left(\frac{q_s}{\sum_{s=1}^{S} q_s}\right)$ is a probability measure over the state space. This probability measure is the risk-neutral or equivalent martingale measure in discrete time. For convenience, we can define $\hat{\pi}_s := \left(\frac{q_s}{\sum_{s=1}^{S} q_s}\right)$ to obtain our final asset pricing formula:

$$p^j = \frac{1}{R^f}\sum_{s=1}^{S}\hat{\pi}_s x_s^j = \frac{1}{R^f}\mathbb{E}^{\mathbb{Q}}\left[x^j\right] \tag{8}$$

Brunnermeier (2015) notes that NA is equivalent to the existence of positive risk-neutral probabilities in the one-period case and notes that a similar theorem exists for the continuous time case.

III.2 Bounding SR

We will derive the Hansen-Jagannathan bound in the one period case. Then, we will use utility theory to bound SR.

III.2.1 The Hansen-Jagannathan Bound

Let the vector of period $(t+1)$ returns to our J assets be given by R_{t+1}. R^f is taken to be fixed across time. We can rewrite (7), first applying it to our portfolio, as:

$$\begin{aligned}
\mathbb{E}\left[h\left(R_{t+1}-R^{f}\right)\right] &= -\frac{\sigma\left(m,hR^{j}\right)}{\mathbb{E}\left[m_{t+1}\right]} \\
&= -\frac{\rho\left(m_{t+1},hR_{t+1}\right)\sigma\left(m_{t+1}\right)\sigma\left(hR_{t+1}\right)}{\mathbb{E}\left[m_{t+1}\right]} \\
\underbrace{\frac{\mathbb{E}\left[hR_{t+1}\right]-hR^{f}}{\sigma\left(hR_{t+1}\right)}}_{\text{Sharpe Ratio}} &= \frac{\overbrace{\rho\left(m_{t+1},hR_{t+1}\right)}^{\text{Pearson correlation}}\sigma\left(m_{t+1}\right)}{\mathbb{E}\left[m_{t+1}\right]}
\end{aligned}$$

The left hand side (LHS hereafter) of the above is the SR of our portfolio. Taking the absolute value of both sides and noting that $|\rho(\cdot,\cdot)| \leq 1$ we obtain the bound for any $h \in \mathbb{R}^{J}$:

$$\begin{aligned}
\left|\frac{\mathbb{E}\left[hR_{t+1}\right]-hR^{f}}{\sigma\left(hR_{t+1}\right)}\right| &= \left|\frac{\rho\left(m_{t+1},hR_{t+1}\right)\sigma\left(m_{t+1}\right)}{\mathbb{E}\left[m_{t+1}\right]}\right| \\
\Rightarrow SR &\underbrace{\leq}_{\text{Hansen-Jagannathan Bound}} \left|\frac{\sigma\left(m_{t+1}\right)}{\mathbb{E}\left[m_{t+1}\right]}\right| = \left|\sigma\left(m_{t+1}\right)R^{f}\right| \qquad (9)
\end{aligned}$$

Most academic work has used realized SRs to obtain lower bounds on the volatility of stochastic discount factors (e.g. the equity premium literature), but in this essay, we will use estimates of upper bounds on $\sigma(m_{t+1})$ to achieve theoretical bounds on achievable SR and thus achievable g.

III.2.2 Utility Theory and the "Good Deal" Bound

The basic intuition behind the GD bound is that utility theory can be used to impose additional restrictions on the SDF; the economics intuition is that our representative agent will purchase any GD assets until the price of those assets is high enough that they are no longer GD assets. Agents' risk-aversion will bound $\sigma(m_{t+1})$ and so it will also bound SR. First, we will derive the SDF as function of the solution to the representative investor's portfolio optimization problem. Once we have done this, we can characterize the bound on $\sigma(m_{t+1})$ by specifying a utility function. We will use a constant-relative risk aversion (CRRA hereafter) because this is both convenient and because CRRA is considered the most realistic utility function within the set of hyperbolic risk aversion utility functions (Brunnermeier, 2015; Danthine and Donaldson, 2005). CRRA utility functions have the following form:

$$u\left(c_t;\gamma\neq 1\right)=\frac{c_t^{1-\gamma}-1}{1-\gamma} \text{ with } u\left(c_t;\gamma=1\right)=\ln\left(c_t\right)$$

where CRRA utility is parameterized by γ and we can obtain the log utility form simply by taking the limit of the LHS case above with $\gamma \rightarrow 1$. Nonetheless, our bound will apply to SR if the marginal investor[5] in

[5] Note that the marginal investor is not necessarily the representative investor. If the representative investor exhibits risk-

the portfolio of interest has a utility function with bounded relative risk-aversion. Relative risk-aversion is defined as:

$$RRA\left(c\right) := -\frac{c_t u''\left(c_t\right)}{u'\left(c_t\right)}$$

In the case of $\gamma = 1$ we have:

$$RRA\left(c_t\right) = -c_t \frac{-\frac{1}{c_t^2}}{\frac{1}{c_t}} = 1$$

We can see that, as one may guess from the name, that RRA is fixed for any CRRA function. The conjecture that the example above suggests, that the RRA of a CRRA function is the γ parameter, is correct. In general, for $\gamma \neq 1$, we can rewrite our CRRA utility as $\frac{c_t^{1-\gamma}}{1-\gamma}$ since this is a monotonic transformation of CRRA. Then, we can see that our RRA is given by:

$$RRA\left(c_t\right) = -c_t \frac{\left(\frac{c_t^{1-\gamma}}{1-\gamma}\right)''}{\left(\frac{c_t^{1-\gamma}}{1-\gamma}\right)'} = -c_t \frac{-\gamma c_t^{-\gamma-1}}{c_t^{-\gamma}} = \gamma$$

We will assume that the our marginal investor is the representative investor throughout the essay and leave the general case of heterogenous investors trading with each other for future work.

In the one-period case our representative's portfolio problem is:

$$\max_h \left\{ u\left(c_0\right) + \mathbb{E}\left[\beta u\left(c_t\right)\right] \text{ such that } \begin{array}{c} c_0 + h \cdot p \leq w_0 \\ \begin{bmatrix} c_1 \\ \vdots \\ c_S \end{bmatrix} \leq h^T X \end{array} \right\} \tag{10}$$

where w_0 is the agent's initial wealth, our expectation is meant to be taken using the objective probability across all states, $\beta \in (0,1)$ is an "impatience" parameter, p is the price vector, c_0 is the initial consumption and c_s is consumption in state s. If all assets are traded, then we can rewrite (10) as:

$$\max_h \left\{ u\left(c_0\right) + \beta \mathbb{E}\left[u\left(c_t\right)\right] | c_0 - w_0 + \sum_{s=1}^{S} q_s \left(c_s - w_s\right) \leq 0 \right\}$$

where we have reduced our constraint to a single dimension. Written this way, we have a canonical microeconomic optimization problem. Because it has a single constraint which will bind in equilibrium (since utility

aversion, then the results in this section will follow.

is increasing in consumption), we can solve the above by the Lagrange method to obtain the following:

$$q_s = \beta\pi \frac{u'(w^s)}{u'(w^0)}$$

so that, recalling the definition of risk-neutral probability and SDF, we have:

$$\frac{\hat{\pi}_s}{\pi_s} = pm_s = \frac{\beta}{R^f} \cdot \frac{u'(w^s)}{u'(w^0)} \text{ and } R^f = \frac{\mathbb{E}[u'(w)]}{u'(w^0)}$$

which can be combined to obtain:

$$\frac{\hat{\pi}_s}{\pi_s} = \frac{u'(w^s)}{\mathbb{E}[u'(w)]} \Rightarrow m_s = \mathbb{E}[m] \frac{u'(w^s)}{\mathbb{E}[u'(w)]} \tag{11}$$

so, that in equilibrium, the utility function and risk-free rate completely specify the SDF:

$$m = \mathbb{E}[m] \frac{u'(w)}{\mathbb{E}[u'(w)]} \tag{12}$$

where we have obtained (12) by applying (11) to find each state's SDF. Intuitively, we can see that if the variance of wealth is finite, $\mathbb{E}[m]$ is well-defined and our utility function is concave, then we should be able to bound the variance of m as well. We can see that every utility function and security structure defines a kernel. Once a complete security structure has been specified, the kernel is uniquely defined in terms of the utility function.

III.2.3 Using the Arrow-Pratt Theorem

Recall that our objective was to find $\sigma(m_{t+1})$; we recast our objective as obtaining $\mathbb{E}\left[m_{t+1}^2\right]$ since we already have $\mathbb{E}[m_{t+1}]$ from our risk-free rate and $\sigma(m_{t+1}) = \sqrt{\mathbb{E}\left[m_{t+1}^2\right] - \mathbb{E}[m_{t+1}]^2}$. Ross (2005) provides the following theorem, which will prove helpful.

First, Ross recalls the Arrow-Pratt theorem which says that if U is a monotone, concave utility function, then $G(U)$ is more risk-averse than U if and only G is a concave, monotonic transform. Then, Ross writes

$V = G(U)$ and applies (12) to V to obtain:

$$\begin{aligned}
\mathbb{E}\left[m_V^2\right] &= \mathbb{E}\left[m\right]^2 \frac{\mathbb{E}\left[v'(w)^2\right]}{\mathbb{E}\left[v'(w)\right]^2} \\
&= \underbrace{\mathbb{E}\left[m\right]^2}_{\text{depends only on } R^f} \frac{\mathbb{E}\left[\left(g\left(u\left(w\right)\right)'\right)^2\right]}{\mathbb{E}\left[g\left(u\left(w\right)\right)'\right]^2} \\
&= \underbrace{\frac{\mathbb{E}\left[\left(g'\left(u\left(w\right)\right)\left(u'\left(w\right)\right)\right)^2\right]}{\mathbb{E}\left[\left(g'\left(u\left(w\right)\right)\left(u'\left(w\right)\right)\right)\right]^2} \frac{\mathbb{E}\left[u'(w)\right]^2}{\mathbb{E}\left[u'(w)^2\right]}}_{\alpha} \underbrace{\mathbb{E}\left[m\right]^2 \frac{\mathbb{E}\left[u'(w)^2\right]}{\mathbb{E}\left[u'(w)\right]^2}}_{\mathbb{E}\left[m_U^2\right]} \\
&= \alpha \mathbb{E}\left[m_U^2\right]
\end{aligned}$$

Now, Ross theorem is true iff $\alpha \geq 1$; this is true if:

$$\frac{\mathbb{E}\left[X^2Y^2\right]\mathbb{E}\left[Y\right]^2}{\mathbb{E}\left[XY\right]^2\mathbb{E}\left[Y^2\right]} \geq 1 \iff \frac{\mathbb{E}\left[X^2Y^2\right]}{\mathbb{E}\left[Y^2\right]} \geq \frac{\mathbb{E}\left[XY\right]^2}{\mathbb{E}\left[Y\right]^2} \tag{13}$$

Ross writes:

$$\mathbb{E}\left[XB\right] \geq \mathbb{E}\left[XC\right] \text{ where } B = \frac{Y^2}{\mathbb{E}\left[Y^2\right]} \text{ and } C = \frac{Y}{\mathbb{E}\left[Y\right]}$$

where one can see this by noting that: $\mathbb{E}\left[B - C\right] = \frac{\mathbb{E}\left[Y^2\right]}{\mathbb{E}\left[Y^2\right]} - \frac{\mathbb{E}\left[Y\right]}{\mathbb{E}\left[Y\right]} = 0$ and that $B - C$ is monotonically increasing with a single root. In our case, X is a positive and increasing function of Y so we have:

$$\mathbb{E}\left[X\left(B-C\right)\right] = \mathbb{E}\left[X\right]\underbrace{\mathbb{E}\left[B-C\right]}_{=0} + \underbrace{\sigma\left(B-C, X\right)}_{>0} \Rightarrow \mathbb{E}\left[X\left(B-C\right)\right] > 0$$

so that we have:

$$0 < \mathbb{E}\left[X\left(B-C\right)\right] = \frac{\mathbb{E}\left[XY^2\right]}{\mathbb{E}\left[Y^2\right]} - \frac{\mathbb{E}\left[XY\right]}{\mathbb{E}\left[Y\right]} \tag{14}$$

To prove (13), Ross defines a new measure by:

$$\mathbb{E}'\left[\cdot\right] = \frac{\mathbb{E}\left[\cdot Y^2\right]}{\mathbb{E}\left[Y^2\right]}$$

and notes that (14) implies:

$$\mathbb{E}'\left[x\right] > \frac{\mathbb{E}\left[xy\right]}{\mathbb{E}\left[y\right]}$$

but then we can simply use Jensen's inequality:

$$\frac{\mathbb{E}\left[X^2Y^2\right]}{\mathbb{E}\left[Y^2\right]} = \mathbb{E}'\left[X^2\right] > \mathbb{E}'\left[X\right]^2 \geq \left(\frac{\mathbb{E}\left[XY\right]}{\mathbb{E}\left[Y\right]}\right)^2$$

which is slightly stronger than what we needed to prove. *QED*.

Now, for any utility function with bounded RRA, we can specify a CRRA function with higher RRA over the entire consumption space. We know that the SDF standard deviation generated by a CRRA function with higher RRA will bound the standard deviation of the true SDF. Below, we derive the bound on g, in the general case. First, recall:

$$\left|\sigma\left(m_{t+1}\right)R^f\right| \geq SR \text{ for any strategy}$$

So we have:

$$R^f\underbrace{\mathbb{E}\left[m\right]}_{\frac{1}{R^f}}\sqrt{\frac{\mathbb{E}\left[u'\left(w\right)^2\right]}{\mathbb{E}\left[u'\left(w\right)\right]^2}-1} \geq \left|\sigma\left(m_{t+1}\right)R^f\right| > \underbrace{\frac{g}{\sqrt{\kappa-g^2}}\overbrace{\sqrt{252}}^{\text{trading days in a year}}}_{\text{annualized market timing } SR}$$

Since the risk-free rate vanishes completely from our LHS, we can simplify our bound and we do so below:

$$\sqrt{\frac{\mathbb{E}\left[u'\left(w\right)^2\right]}{\mathbb{E}\left[u'\left(w\right)\right]^2}-1} > \frac{g}{\sqrt{\kappa-g^2}}\sqrt{252} \tag{15}$$

Now, we can rearrange the terms in (15) to obtain (16), which is our final bound:

$$\sqrt{\frac{\left(\frac{\mathbb{E}\left[u'(w)^2\right]}{\mathbb{E}\left[u'(w)\right]^2}-1\right)\kappa}{\frac{\mathbb{E}\left[u'(w)^2\right]}{\mathbb{E}\left[u'(w)\right]^2}-251}} > g \tag{16}$$

We can see that we only need to know the upper bound on RRA for the marginal investor in our economy, our marginal investor's end of period distribution of wealth w and low-order statistics about the asset itself to bound market timing ability g.

IV Too Good to Be True: A Bit of Empirical Work

Now, we want to show that the bound we have constructed in (16) has some bite by applying it to one of the studies shown in Table 1.[6] First, we must choose an upper bound on RRA. Choosing a higher RRA gives

[6] In fact, all of the studies shown in Table 1 (for which data was available) have been checked. With the exception of Wuthrich

a looser bound because the LHS of (16) is monotonically increasing in RRA. Second, we will specify a final distribution for wealth. There is no theoretical basis for making either choice and so our discussion on RRA upper bound selection draws on a combination of introspection and empirical work on individual preferences while we simply choose a common specification (the log-normal) for the final distribution of wealth.

IV.1 Justification for Parameter Choices and Functional Specifications

Potì and Wang (2010) provide a brief summary of the literature on upper bounds for RRA which we review. Ross (2005) notes that an RRA greater than five for the marginal investor "would imply that the investor is willing to pay more than 10% per annum to avoid 20% volatility"; since the S&P 500 has approximately this volatility and mean, and must be held in equilibrium by the marginal investor[7], under CRRA we can reject RRA greater than five. We note that this bound is lower than the upper bound on RRA discovered by Friend and Blume (1975); Meyer and Meyer (2005), who look at household asset allocation data to estimate RRA, and find a range of 2.0 to 6.4. On the other hand, Potì and Wang (2010) point out that the data presented by Barsky et al. (1997) implies an RRA between 0.8 and 1.6. We will choose an RRA of 6.5, as no study of individual risk-tolerance we have seen has ever discovered any individual with an RRA higher than 6.5.

IV.1.1 Some Subtleties

We will note that we can always convert a (g, κ) pair into SR and that a violation of the upper bound on g is also a violation of the upper bound on SR. Our empirical exercise relies crucially on (what economists consider to be) a subtle difference between ex-ante SRs, which must be estimated from statistical models and ex-post SRs, which are reported in the studies we assess.[8] For example, a portfolio consisting of two lottery tickets, one of which (unknown to the purchaser) is a winner, would have a realized SR exceeding the bound, but the purchaser's rational ex-ante return expectation would be very low and so would not violate our NA bound. Fortunately, if the results presented in Table 1 are not simply Type I errors, then they should be subject to the bounds we constructed in III. The first use of the SR in Sharpe (1966) finds not a single mutual fund able to achieve an SR of 1.0. Shen (2002), describes one market timing strategy that "actually worked" as having achieved an SR of 0.73. None of the strategies he simulates produce an SR greater than 0.8. All of the SR bounds implied by (16) are greater than 1 and so in this sense our bound is

et al. (1998), all violate the GD bound we construct in this section.

[7] Otherwise its price would decrease until it was held by the marginal investor.

[8] In practice, GD results must lead to both predicted and realized exceedance of the SR bound. This de facto criterion addresses the joint-hypothesis problem: a flawed statistical model could lead to many assets ex-ante exceeding the SR bound, but it is difficult to argue that the model is flawed if those assets subsequently realize SRs exceeding the bound.

a priori reasonable.

IV.1.2 Wealth Shocks

Like Ross (2005), we will assume that the individual's end of period wealth is the product of multiplicative wealth shocks from trading the asset of interest and so it is log-normally distributed. In other words, we have:

$$\ln(w) \equiv N\left(\mu, \sigma^2\right) \tag{17}$$

We might obtain this by assuming that our market contains a single log-normally distributed asset of interest and a risk-free asset; if we work with the discrete state security structure proposed in III we can choose state payoffs and probabilities for a single risky-asset economy such that the distribution of the end of period wealth becomes arbitrarily close to a log-normal distributions.

IV.2 Evaluating a Study

We will show how the bound can be used to assess the economic signficance of a reported finding using Bollen et al. (2011) and Vu et al. (2012). We chose Bollen et al. (2011) and Vu et al. (2012) for exhibition in this essay for three reasons.

1. The Dow Jones Industrial Average (DJIA) and stocks of Apple, Google, Microsoft and Amazon are liquid assets and there is reliable data available; the results reported in this essay are easily verified.

2. With the exceptions of Huang et al. (2010), the algorithm used by the authors has the highest g reported.

3. Shockingly, although given the extremely high predictive power they report, unsurprising, Vu et al. (2012) appear to use their training set as their testing set. Given the obvious bias this induces in their reported g, our procedure's ability to finger these predictability results as likely Type I errors can be characterized as a a sanity check. Similarly, Pav (2012) suggests that the widely-cited results reported in Bollen et al. (2011) are the result of computational and statistical errors.

First, we calculate $\mathbb{E}\left[m^2\right]$ under CRRA utility $\frac{c^{1-\gamma}}{1-\gamma}$ with $\gamma = 6.5$. Using the moment generating function to write $\mathbb{E}\left[w^{-R}\right] = \exp\left[\frac{1}{2}(\gamma\sigma)^2 - \gamma\mu\right]$, we can obtain:

Asset	Reported g	$\hat{\kappa}$	$\hat{\sigma}$	Upper Bound of g	Violated
DJIA	36.67%	0.61	0.26	20.12%	Yes
AAPL	32.93%	0.37	0.18	6.72%	Yes
GOOG	30.49%	0.72	0.27	23.57%	Yes
MSFT	25.61%	0.42	0.17	6.54%	Yes
AMZN	25.00%	0.52	0.30	26.87%	No

Table 3: This table contains the maximal accuracy results for predicting the up-down pattern of stock price moves reported in Bollen et al. (2011) and Vu et al. (2012). Notice that it appears Nassirtoussi et al. (2014) misreported their maximal accuracy as 84.62% when in reality the maximal accuracy reported was 82.93%. However, Nassirtoussi et al. (2014) make reference to "online tests" (i.e. ex-sample tests) but no tests are identified as online or ex-sample in Vu et al. (2012); it is possible that online tests referred to by Nassirtoussi et al. (2014) were presented at the conference but not reported in the final, published proceedings. For Vu et al. (2012) statistics were calculated over 41 trading days from April 1, 2011 to May 31, 2011 while for Bollen et al. (2011) we use all trading days from December 1, 2008 to December 19, 2008. Single-name equity data is obtained from Google Finance while the DJIA data is obtained from the St. Louis Federal Reserve FRED database.

$$\begin{aligned} \mathbb{E}\left[m^2\right] &= \frac{\mathbb{E}\left[w^{-2\gamma}\right]}{\mathbb{E}\left[w^{-\gamma}\right]^2} \\ &= \frac{\exp\left[2\left(\gamma\sigma\right)^2 - 2\gamma\sigma\right]}{\exp\left[\left(\gamma\sigma\right)^2 - 2\gamma\sigma\right]\left(R^f\right)^2} \\ \Rightarrow \sqrt{\exp\left\{\left(\gamma\sigma\right)^2\right\} - 1} &> \left|\sigma\left(m\right)R^f\right| \end{aligned}$$

Now, we can use the fact that our LHS bounds SR for any portfolio and repeat the algebra leading to (16) to write out the inequality:

$$\sqrt{\frac{\kappa\left(\exp\left\{\left(\gamma\sigma\right)^2\right\} - 1\right)}{251 + \exp\left\{\left(\gamma\sigma\right)^2\right\}}} > g$$

We can estimate $\hat{\sigma} = \sqrt{252}\sqrt{\frac{\sum_{t=1}^{T}(r_t-\mu)^2}{T-1}}$ by obtaining the daily sequence of r_t and μ for each of the four stocks. Since all units are annualized, we must annualize our daily standard deviation by multiplying by $\sqrt{252}$. The true σ of the portfolio, because of the possibility of diversification, must be smaller than this quantity and so by choosing higher values of $\hat{\sigma}$ (derived from a single asset rather than the optimal portfolio) we bias our test towards acceptance of the reported g. Empirically, in Table 3 we evaluate:

$$\sqrt{\frac{\hat{\kappa}\left(\exp\left\{\left(6.5\hat{\sigma}\right)^2\right\} - 1\right)}{251 + \exp\left\{\left(6.5\hat{\sigma}\right)^2\right\}}} > \hat{g} \tag{18}$$

where the LHS is our upper bound and $\hat{g}$ is the predictive ability reported in the study. Wherever the inequality in (18) fails to hold, we have identified either a Type I error or a violation in one of the assumptions driving our analysis. On balance, we strongly suspect that in the case ofBollen et al. (2011) and Vu et al. (2012) we have identified several Type I errors.

IV.2.1 Results, Discussion and Future Work

As we suggested earlier, the upper bounds given by (18) are plausible GD bounds. Mathematically however, the LHS of (18) can take on any value in the range $(0, \infty)$. Therefore, any restrictions that (18) imposes on g arise from the agent characteristics and equilibrium asset price behavior. In other words, (18) is a "no good deal" equilibrium condition, analagous to an NA equilibrium condition. Our literature review suggests that (18) is the first ever GD bound on market timing ability. While it is numerically possible for the LHS of (18) to exceed 0.5 (given that we have chosen upper bounds at every juncture that, collectively, may not be consistent with an actual economic equilibrium) and thus place no restrictions on the data, all of the upper bounds we calculate are between 0 and 0.30. At the same time, because Vu et al. (2012) estimates and tests their model in the same sample, the results reported therein are likely to be Type I errors. Our method's ability to identify these results as such marks it as a promising approach for ex-post evaluation of purported market timing ability. Our method's suggestion of Bollen et al. (2011) suggests that the results therein are worth revisiting.

There are three major families of extensions that would improve the model. These three families of extensions are related to the market timing strategy proposed in II, the one-period model displayed in III.1 and the GD bounds developed in III.2, respectively.

With regards to the first set of extensions, it is difficult to give abstract representation to market timing schemes that predict only raw, but not excess returns and depart from unit leverage. While the current zero interest rate environment makes the former extension only of theoretical interest in the United States at the present moment, this extension is necessary before the technique can be applied to markets in other countries in which the risk-free rate is far from zero. The latter extension is a necessity for realism: virtually all proprietary traders have leverage that co-moves with the business cycle. Including variable leverage should increase the SR of the market timing scheme and decrease the bound on g, while accounting for non-zero risk-free rates should decrease the SR and increase the bound on g.

Modeling the stock market as a set of assets with fixed cash flows independent across time seems like an obvious weakness of our approach. Unfortunately, the full multi-period "Arrow"-style discrete time, discrete state model is complex. At the same time, we would also like to account for incomplete markets. Incomplete markets are likely to lead to looser bounds on g since agents will optimize over different SDFs and Ross (2005) shows that agents prefer more variable SDFs. Obtaining the Hansen-Jagannathan bounds in the incomplete market, discrete time, multi-period, discrete (or continuous) state security structure is a non-trivial, but solved problem and so incorporating this more general structure into our analysis, up to and including the derivation of the Hansen-Jagannathan bound, should be straightforward. A more general security structure may also allow a more principled choice for end-of-period wealth than (17), which was chosen as a modeling convenience.

On the other hand, as far as I know, the GD bounds used in this essay have only ever been developed in the discrete time, discrete state, one-period model. Ross (2005) suggests that the development of GD bounds under the full suite of available security structures is an open problem and so we look forward to future work in this area. Finally, there exist pathological utility functions with unbounded RRA. If one of these utility functions describes the actual behavior of the marginal investor, then all of the analysis above is moot. Ross (2005) suggests that we should be able to discover GD bounds for all risk-averse utilities over lotteries with bounded wealth: giving closed form representation to GD bounds derived from utility functions with unbounded RRA is an active research area today and it is possible that a simple closed form GD bound that applies to all risk-averse utility functions may be discovered. It is likely that extensions on the GD approach will lead to tighter bounds on g and ultimately a more powerful system for detecting Type I errors in market timing ability.

Nomenclature

CRRA constant relative risk aversion

GD good deal

LHS left hand side

NA no arbitrage

SDF stochastic discount factor

SR Sharpe Ratio

References

Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro, "PREFERENCE PARAMETERS AND BEHAVIORAL AN EXPERIMENTAL APPROACH IN THE HEALTH AND RETIREMENT STUDY," *Quarterly Journal of Economics*, 1997, (May), 537–579.

Bollen, Johan, Huina Mao, and Xiaojun Zeng, "Twitter mood predicts the stock market," *Journal of Computational Science*, March 2011, *2* (1), 1–8.

Brunnermeier, Markus K, *Asset pricing I: Pricing Models* 2015.

Danthine, Jean-Pierre and John Donaldson, *Intermediate Financial Theory*, 2 ed., Burlington: Elsevier, 2005.

Friend, Irwin and Marshall Blume, "The Demand for Risky Assets," *American Economic Review*, 1975, *65* (5), 900–922.

Hagenau, Michael, Michael Liebmann, and Dirk Neumann, "Automated news reading: Stock price prediction based on financial news using context-capturing features," *Decision Support Systems*, June 2013, *55* (3), 685–697.

Huang, Shian Chang, Pei Ju Chuang, Cheng Feng Wu, and Hiuen Jiun Lai, "Chaos-based support vector regressions for exchange rate forecasting," 2010.

Mahajan, Anuj, Lipika Dey, and Sk Mirajul Haque, "Mining financial news for major events and their impacts on the market," *Proceedings - 2008 IEEE/WIC/ACM International Conference on Web Intelligence, WI 2008*, 2008, pp. 423–426.

Meyer, Donald J and Jack Meyer, "Relative risk aversion: What do we know?," *The Journal of Risk and Uncertainty*, 2005, *31* (3), 243–262.

Nassirtoussi, Arman Khadjeh, Teh Ying Wah, Saeed Reza Aghabozorgi, and David Ngo Chek Ling, "Text Mining for Market Prediction: A Systematic Review," *Expert Systems with Applications*, June 2014, *41* (16), 7653–7670.

Pav, Steven, "Converting Timing Edge to Sharpe," 2012.

Potì, Valerio and DengLi Wang, "The coskewness puzzle," *Journal of Banking and Finance*, 2010, *34* (8), 1827–1838.

Rachlin, Gil, Mark Last, Dima Alberg, and Abraham Kandel, "ADMIRAL: A data mining based financial trading system," *Proceedings of the 2007 IEEE Symposium on Computational Intelligence and Data Mining, CIDM 2007*, 2007, (Cidm), 720–725.

Ross, Steven, *Neoclassical Finance*, Princeton University Press, 2005.

Schumaker, Robert P, Yulei Zhang, Chun neng Huang, and Hsinchun Chen, "Evaluating sentiment in Financial News Articles," *Decision Support Systems*, 2012, *53* (3), 458–464.

Sharpe, W.F., "Mutual Fund Performance," *The Journal of Business*, 1966, *39* (1), 119–138.

Shen, Pu, "Market Timing Strategies that Worked," 2002.

Soni, Ankit, Nees Jan van Eck, and Uzay Kaymak, "Prediction of Stock Price Movements Based on Concept Map Information," in "2007 IEEE Symposium on Computational Intelligence in Multi-Criteria Decision-Making" IEEE 2007, pp. 205–211.

Vu, Tien-Thanh, Shu Chang, Quang Thuy Ha, and Nigel Collier, "An experiment in integrating sentiment features for tech stock prediction in Twitter," in "Workshop on Information Extraction and Entity Analytics on Social Media Data" 2012, pp. 23–38.

Wuthrich, B., V. Cho, S. Leung, D. Permunetilleke, K. Sankaran, and J. Zhang, "Daily stock market forecast from textual web data," *SMC'98 Conference Proceedings. 1998 IEEE International Conference on Systems, Man, and Cybernetics (Cat. No.98CH36218)*, 1998, *3*, 2720–2725.

Zhai, Yuzheng, Arthur Hsu, and Saman K Halgamuge, "Combining News and Technical Indicators in Daily Stock Price Trends Prediction," *Advances in Neural Networks – ISNN 2007*, 2007, pp. 1087–1096.