Feature and expert-based market prediction using online low-regret algorithms

By Chris Murray and Geoff Gordon

Abstract

For decades, people have expended untold amounts of effort trying to predict the stock market. We frame market prediction as an experts problem, in the online learning sense, and show that simple strategies have optimal regret properties. We show how to generalize these strategies to more exotic types of portfolios. We also show how to take advantage of recent empirical results in finance by creating feature portfolios, and how these can be used to achieve dramatic returns within the low-regret framework.

1. Introduction[1]

An abundant literature exists relating to predicting movements in the stock market. From snake-oil salesmen and questionable websites selling advice with completely unverifiable value to seminal papers in top peer-reviewed journals, a tremendous amount of effort has been put into understanding stock market returns. This research will draw from the latter type of sources. One goal will be to gain a deeper understanding of the robustness of previously reported results and see how observable features of a stock's history influence its future return.

One way or another, we will go about developing procedures for trading stocks (or other assets) in the market. In our view, there are two ways to go about this. The first is to develop a ***heuristic trading algorithm***.[2] By heuristic, we mean algorithms which make some assumptions about how the market works and try to profit from these assumptions. These processes can be based on simple intuition or on arbitrarily complex statistical learning techniques, but either way they will depend on assumptions about how the market evolves and may perform poorly.

The second type of procedure is what we call a ***performance bounded trading algorithm***. This is a trading process which has a powerful worst-case performance guarantee that will hold regardless of how the market evolves. We will develop a trading process that is performance bounded, rather than heuristic, and the next section will try to motivate our reasoning. These performance bounded algorithms are partly theoretical and partly practical: we can theoretically prove that the algorithms will have low regret, but empirically this is only useful if the best asset in the algorithm's universe performs relatively well. We will frame investment as an online learning experts problem, and will give performance guarantees with respect to the best expert chosen in hindsight. We will show how to integrate "features" of various stocks into this process, in order to make bets on how stocks with different features will perform over time. We will choose features that are broadly applicable to all stocks at all times: this excludes event-type features like a recent IPO, dividend omission, acquisition, or rare technical indicator. We will also only use features which have been identified in peer-reviewed journals as relating to a stock's future returns and having some possible economic meaning. This is partly done to avoid data mining biases[3] and means that features such as "The stock's 50 day moving average crossed above its 200 day moving average and below its 1-year moving average between 18 and 23 days ago on medium-high volume" will be excluded.

---

[1] This paper would not have been possible without the help of my advisor, Geoff Gordon.

[2] Heuristic does not in any way imply that these processes are less sophisticated, less mathematically formal, or in any way inferior. It simply means that they do not give performance guarantees.

[3] Though, to some extent, it simply means that we have outsourced our data mining.

1.1. Heuristic versus performance bounded trading processes

I will contrast heuristic with performance bounded trading processes by example. I will start by describing in detail two heuristic trading algorithms and two performance bounded algorithms from recent literature. Two recent papers on heuristic trading processes, both by machine learning experts, are "Can We Learn to Beat the Best Stock" (Borodin, El-Yaniv, and Gogan 2003,2004) and "Automated Trading with boosting and expert weighting" (Creamer and Freund 2006). The performance bounded algorithms I will look at are the classic "Universal Portfolio" algorithm (Cover 1991) and Online Newton Step (Agarwal et. al. 2006)

The Borodin et. al. (2003, 2004) papers develop ANTICOR, an algorithm for trading based on various types of correlation. It uses one parameter, **w**, a window length of past trading days. At any time t, the algorithm considers the past two windows **w2** and **w1**: **w2** is [**t-2w+1** to **t-w**] and **w1** is [**t-w+1** to **t**]. It defines M(i, j) as the correlation of stock i's log returns over **w2** with stock j's log returns over **w1:** that is, how much stock i's return in the past window correlates with stock j's return in the next window. Also, it defines R(i) and R(j) as the mean log return over the most recent window **w1** for stocks i and j. The algorithm then decides on a way to allocate wealth among stocks over the next trading day, **t+1**. To do this, it defines a variable **Claim(i→j)**, which is initialized to zero. If R(i) > R(j) and M(i,j) > 0, then

Claim(i→j) = M(i,j) + MAX(0, -M(i,i)) + MAX(0,-M(j,j))

After defining all these **Claim(i→j)** variables, the algorithm adjusts its current portfolio by transferring wealth from stock i to stock j when **Claim(i→j)** is positive. The authors also describe a sort of $ANTICOR^2$ algorithm which uses ANTICOR to choose an allocation among other ANTICOR algorithms with different window sizes. Once we get through the notation, we see that ANTICOR is making bets on the consistency of measured ***positive lagged cross-correlation*** and ***negative autocorrelation***.[4] The authors claim that ANTICOR earns outrageous returns on most of the data sets they use (like 100% per year).

The next heuristic algorithm, automated trading with boosting and expert weighting (Creamer and Freund 2006) builds a three-layer trading system, based on sophisticated machine learning tools. The algorithm uses pre-defined features, which are a large range of 30 somewhat exotic sounding technical indicators, which are often complex ratios of moving averages, stochastic trading rules, or other indicators with questionable economic meaning. The algorithm's first layer consists of a set of alternating decision trees (ADTs) trained on the features to predict return. The next layer uses boosting and expert weighting to choose among ADTs. The last layer, described as a risk management layer, prevents trading on weak signals. They claim this system earns abnormal returns of roughly 6% to 20% out of sample, depending on simulated transaction costs.

[4] Autocorrelation is a time-series' correlation with its ***own*** past values. Lagged cross-correlation is one time series' correlation with another time series' past values.

The performance bounded algorithms are, by contrast, much simpler. However, since they have performance guarantees, some non-trivial amount of work is involved in showing that the guarantees hold. The universal portfolio algorithm (Cover 1991, Blum and Kalai 1997) develops a method to achieve the same long-term growth rate as the best constantly rebalanced portfolio (CRP) chosen in hindsight. This is a stronger guarantee than the best individual stock chosen in hindsight. Details are given later, but it basically allocates wealth initially to every CRP on the simplex, using a uniform or Dirichlet prior, and from there simply buys and holds each CRP. The log performance of this algorithm will lag the best constantly rebalanced portfolio by a constant of roughly the number of experts times the log of the number of periods.

Online Newton Step (Agarwal et. al. 2006) is a newer algorithm for the portfolio selection problem. Its allocation every period is based on a second-order approximation to the current performance of each expert, but some fraction of wealth is allocated uniformly equal-weight every period. This has performance bounds that are roughly like those of Universal Portfolio.

Contrasting the heuristic with the performance bounded algorithms, several differences are clear. We will use these differences to try to justify our use of performance bounded, rather than heuristic, algorithms.

1. Heuristic algorithms are making reasonably complex bets, while performance bounded algorithms aren't.

The ANTICOR algorithm makes a bet based on stability of cross-correlation and autocorrelation. Creamer and Freund's algorithm makes very complex black-box-like stationarity bets that we can't hope to understand well. That is, given a market sequence, it would be hard to predict how well their algorithm would perform on that sequence without running the algorithm. By contrast, the performance bounded algorithms are both relatively bland: universal portfolio is a simple application of buy-and-hold. We argue that it's OK to make complex, even black-box, bets: that's probably the only way to make a lot of money very fast. However, no matter how sophisticated we are about it, it's still gambling, and adding sophistication often makes it harder for us to understand precisely the bets we are making.

2. Heuristic algorithms sound cooler and can use more fancy tools.

Let's face it: the Creamer and Fruend algorithm sounds a lot more exciting than Universal Portfolio. It puts together ADTs, a somewhat fancy kind of decision tree. Then it uses relatively recent tools like boosting and expert weighting to further train and choose among these ADTs, with a risk management layer on top of that. That's a lot neater than a simple buy-and-hold strategy like Universal Portfolio. It's ***fun*** thinking about how we could combine the machine learning tools we know and love to make a great stock picking algorithm.

However, we can run into some trouble this way. As we try different tools to see what does and doesn't work, we will almost surely over-fit our training and test data unless we are ***exceptionally*** careful to keep a true holdout set that we don't touch until we have selected our algorithm and all its parameters. Even if researchers all do this, published results will still be biased if many researchers find but don't publish null results, while the few researchers who do find positive results submit publications. Over-fitting is an unusually important problem in stock market data because stock returns are, by nature, very noisy and driven by covariates that we can't condition on (such as news reports, investor optimism during a particular month, footnotes in accounting statements, and even noise on the trading floor) and because there are a limited number of stocks and only one global economy. Any time we fit a regression to individual stock data and get an R^2 of 0.7, or we find a hyperplane that separates all stocks with positive return from those with negative return, we are shamefully guilty of over-fitting and need to go back to square one.

Also, and perhaps more importantly, as we use more and more complex heuristics or combinations of machine learning tools to predict the market, any results we obtain may not generalize well to other problems or data sets.

3. Heuristic algorithms claim higher returns than performance bounded algorithms

Heuristic algorithms are, in a sense, optimized to do well on the data on which they are run (or closely related data). Performance bounded algorithms must be more conservative since they have to do "well" on the worst possible data set in existence. On all the datasets I've looked at, the performance-bounded algorithms do not earn exceptionally high returns.

4. Heuristic algorithms aren't robust.

Heuristic algorithms may perform extremely well under some market conditions but extremely poorly under others. Performance guaranteed algorithms will perform well, by some measure, no matter what.

Thus despite the initial appeal of trying to develop heuristic algorithms which earn exceptionally high returns, we feel that it is more important to look at techniques which have performance guarantees.

2. Market Theory and Empirical Results

This section introduces the reader to portfolio theory and the empirical results from the finance literature that we use to construct our market features. It may be skipped without lack of continuity.

2.1. The Evolution of Efficient Market Theory

The portfolio ideas of Markowitz (1959) and the Sharpe (1964), Lintner (1965), and Black (1972) model had for a long time convinced academics and, to a lesser extent practitioners, that a stock's expected return (over the risk-free rate) should primarily depend linearly on its covariance with the value-weighted market portfolio (beta) and nothing else. This, roughly, is often called the CAPM (capital asset pricing model). The CAPM states

$$E(R_i) = \alpha + R_f + \beta_i \times E(R_M - R_F) \qquad \textbf{(2.1)}$$

where

$$\beta_i = \frac{\mathrm{Cov}(R_i, R_M)}{\mathrm{VAR}(R_M)} \qquad \textbf{(2.2)}$$

α is a measure of excess return that is zero in equilibrium, $\mathbf{R_i}$ is the return on asset **i**, $\mathbf{R_m}$ is the market return, and $\mathbf{R_f}$ is the risk-free rate (often taken to be the yield on a short-term treasury bill). Note that this return relation in the CAPM is not an assumption: it is a conclusion that follows from rational risk-averse investors. Qualitatively, in a market where information about the future distribution of returns and covariances is publicly known, a risk-free rate with unrestricted borrowing and lending exists, there is only one time period to invest, and investors are risk-averse, the CAPM is an equilibrium result. A derivation can be found in Norstad (2005), which is based on investors optimizing a utility function that has the form $\mathbf{U = E(R) - \delta \times Var(R)}$. In this setting, equation (2.1) must hold if investors are properly optimizing their utility.

Despite its simple and appealing nature, the CAPM has generally been a complete failure empirically at explaining the cross-section of stock returns. That is, if we look at historical data, we don't find a strong relation between a stock's sample covariance with the market and its future return.

In the 1980s even academics' view of the market began to change as studies found that the cross-section of stock returns is related to companies' market capitalization, earnings to price, book to market value, cash flow to price, and past sales growth (Banz 1981, Basu 1983, Rosenberg, Reid, and Lanstein 1985, and Lakonishok, Shleifer, and Vishny 1994). Fama and French (1992,1996) found that a three-factor linear model which relates the cross-section of returns to covariance with the market, a factor related to the ratio of a stock's book value to its market value, and a factor related to firm size does a good job of explaining returns of portfolios sorted on Cash Flow to Price, sales growth, book to market, earnings to price, etc. Fama and French also argue that their three factors may proxy for some market-wide distress-related risk which investors should be compensated for holding in a multi-factor world (Cochrane 1999 and references therein). The Fama-French model is widely used and cited today. Another anomaly which has largely defied explanation is medium-term return persistence (Jegadeesh and Titman 1993).

2.2. Efficient Market Theory and Multifactor models

One point often made by Eugene Fama, one of the most well respected academics in finance, is that it's impossible to directly do a hypothesis test regarding market efficiency: we can only do a test of the ***joint*** hypothesis that the market is efficient AND our return model is correct. Rejection may mean that the market is inefficient, or that the model we are using to specify “correct” equilibrium returns is wrong. Multifactor models have multiple “priced” sources of systematic risk for which investors can be compensated (see Cochrane 1999a, 1999b). For example, investors might want to avoid both market risk, for which beta is a proxy, and an extra risk factor relating to stocks which do poorly in recessions, since the average investor would not want his or her portfolio to do poorly in a recession. There seems to be a consensus in the academic community that multifactor models are a more correct view of the world, though there isn’t agreement on which factors should be included.

Empirically, a stock’s $\boldsymbol{\beta}$ isn’t priced: that is, stocks with higher $\boldsymbol{\beta}$ don’t tend to have higher return. If this was the case in a world where we knew that investors care only about mean and variance, we would argue that the market is inefficient since stocks with higher $\boldsymbol{\beta}$ risk don’t provide enough return to compensate for that risk. However, in a mutlifactor world we can’t make such a statement, since stocks can be riskier or less risky in some way that $\boldsymbol{\beta}$ doesn’t capture, so the market may still be efficient even if $\boldsymbol{\beta}$ isn’t related to future returns. There are disagreements on what constitutes a systematic risk factor. For example, Lakonishok, Shleifer, and Vishny (1994) find that stocks with value properties have higher returns: they perform a detailed empirical analysis to argue that these stocks aren’t any riskier. Fama and French (1993) argue that these value properties may proxy for some recession-related risk factor.

One thing to note is that only systematic risk factors should be priced, since non-systematic ones can be diversified away. Imagine this wasn’t the case, and a portfolio’s expected return was related to its variance. Then we could argue that a portfolio of stocks belonging to any one industry is more risky than the market as a whole (which is correct), and thus that industry portfolios should have higher expected return than the market to compensate for this risk. But the latter conclusion is impossible because the market is simply a weighted average of industry portfolios. Even systematic risk factors may be very hard to identify by only looking at returns. A strategy that writes (sells) out-of-the-money put options generates returns that seem attractive over five or ten-year periods, but bears tremendous systematic risk (see Lo 2001).

2.3. Return Decomposition in the CAPM

The reason we covered theoretical finance was to be able to decompose returns of a portfolio. When we invest, we are able to decompose the returns of any asset into three components:

1) We earn a risk-free rate $\mathbf{R_f}$: this is what market participants must be given to convince them to part with their money. This is like the return on a money market account or on a treasury bill.
2) We earn a risk-premium for investing in assets with systematic risk.

3) We may earn alpha, excess return that is over and above compensation for risk and that comes about through our investment skill or luck. In most models, alpha must be zero-sum.

The CAPM return decomposition is given in equation 2.1. It's important to remember that **β** and **α** are ***relative*** measures of risk and excess return. Just because an asset has positive alpha doesn't mean it's an attractive asset to hold in isolation. It means that, when added to the market portfolio, that asset will allow an investor to earn higher return for a given level of variance.

2.4. Related empirical work

We use empirical finance research to derive our features. These features roughly fall into four categories. The first deals with observed medium-term momentum in stock returns. That is, stocks that did well in the medium-term past often do well in the medium-term future. Roughly, there are two kinds of momentum: earnings momentum and return momentum. Earnings momentum strategies buy stocks with recent strong earnings or earnings surprises, while return momentum strategies buy stocks with strong past returns. Neither strategy subsumes the other, and stocks with strong earnings momentum still perform well controlling for past returns, and vice versa (Chan, Jegadeesh, and Lakonishok 1996). However, earnings momentum is more difficult to define objectively, and unless implemented with great care may suffer from look-ahead bias[5] and survivorship bias since historical earnings data are less complete than historical return data. Thus we will use the term "momentum" strategies to refer to return momentum. The seminal study of this was by Jegadeesh and Titman (1993) who found excess returns of about 1% or more per month to the following strategy (for all K and L).

1) Every month sort stocks by their returns over the past (K=3,6,9,12) months, put the stocks into equal-weight deciles, and go long the decile with the highest past return and short the decile with the lowest past return. Hold this position for (L=3,6,9,12) months. Every month apply (1/L) of your capital to this strategy.

Other studies have found that this return persistence is an international phenomenon and international momentum returns correlate with those in the U.S., suggesting a common factor (Rouwenhorst, K Geert. 1998). Also, country indexes themselves show return momentum (Asness, Liew, and Stevens 1997). There is evidence that momentum in industry indexes explains much of individual stocks' momentum (Moskowitz and Grinblatt 1999), and that a stock's proximity to its 52-week high is a more significant predictor than return momentum and explains much of it, possibly suggesting a behavioral cause (George and Hwang 2004). Other authors think transaction costs will eliminate any momentum profits (Lesmond, Schill, and Zhou 2004, Keim 2003) due to frequent trading in illiquid

---

[5] Look-ahead bias occurs when a strategy uses data that wouldn't be available at that time. For example, in the U.S. fiscal year earnings may not be released until four months after the year ends.

stocks. There has been no satisfying explanation of momentum profits. Conrad and Kaul (1998) point out that, in a simple model where every stock has constant expected return, past returns are a noisy proxy for expected return, and momentum strategies will roughly buy high-expected return stocks and sell low-expected return stocks even if there is no time-series predictability in security returns. In this model momentum profits increase with cross-sectional variation in expected returns. However, momentum-based portfolios show a reversal after 12 months, which violates this model, and Jegadeesh and Titman (2002) argue that the methods used by Conrad and Kaul to estimate the cross-sectional variation in expected returns were flawed.[6] Hong and Stein (1999) present a model where two types of rational agents, news watchers and momentum traders, trade in such a way as to generate momentum profits followed by longer-term reversals.

The second feature we use is a measure of a stock's "value" properties. Since the price of a stock is the discounted expected value of its future cash flows, stocks with high projected future cash flow growth tend to have high prices relative to their current "book" values[7]. Many financial ratios are proxies for value properties, including earnings to price (E/P), cash flow to price (C/P), dividend yield (D/P), and book value to market value (B/M). Lakonishok, Shleifer, and Vishny (1994) find that stocks with value characteristics (having high values of the above four variables) have above-average performance for several years, and claim that these stocks are not riskier but rather that their excess return is due to sub-optimal investor behavior.

Fama and French (1992,1993,and 1996) claim that the above four values are largely redundant, and they use book to market value to measure how much a given stock is a "value" stock. They created a three-factor model to explain stock returns based, roughly, on a stock's market capitalization (size), the extent to which it's a "value" stock, and its beta. They argue that value stocks and smaller capitalization stocks earn a premium due to their higher risk: thus to understand if a stock or portfolio's return is abnormally high, they adjust its return for market risk (β), "value stock" risk, and small capitalization stock (liquidity) risk. Fama and French claim that value stocks may demand a premium in an efficient market because value is a proxy for undiversifiable distress risk. Smaller capitalization stocks may carry a premium because they are illiquid and there is less information about them (see Merton (1987) for an example of how lack of information can lead to a return premium).

Stock momentum, however, is not explained by the Fama-French three-factor model. Asness (1997) finds that value characteristics are negatively

---

[6] Cross-sectional variation in expected returns, on a given date, is the variance, across stocks, of underlying expected returns. However, expected return is not observed. Conrad and Kaul used a stock's past realized return as a proxy for its expected return. But if a stock's realized return is its expected return plus an error term, than the cross-sectional variance of realized returns is the cross-sectional variance of expected returns ***plus*** the variance of the error terms. Remembering that in the market error terms are large, we see that this can bias results significantly.

[7] Book value is the company's assets minus its liabilities. It VERY ROUGHLY says what the company's value is without future earnings.

correlated with momentum characteristics, though both are positively related to future return. He also finds that conditioning on a stock's value properties and its momentum properties jointly can explain returns better. Arshanapalli, Coggin, and Doukas (1998) find value stocks perform better than growth stocks throughout the developed world. Asness, Friedman, Krail, and Liew (2000) show that the return spread between value and growth can be forecast, and they gained major credibility by publishing in Oct 1999 the fact that their model was forecasting a 52% return spread for value over growth (following their forecast, many internet stocks crashed, almost none of which were "value" stocks).

The third major feature is firm size. Small stocks have historically outperformed larger ones (Banz 1981, Fama and French 1992), though that premium has decreased recently and largely occurs in January (Roll 1983, Branch and Chang 1985) and may be tax related. As mentioned, firm size is one constituent of the Fama-French model (1993, 1996) and it plays a part in momentum strategies since stocks with extreme past performance tend to be smaller (Jegadeesh and Titman 1993). More subtly, many of the above studies use results from equal-weight portfolios (as does this paper) to justify their claims (Jegadeesh and Titman 1993 among others). Equal weight portfolios under-weight large stocks and overweight small stocks, and can cause results to be driven by a few small stocks that make up only a tiny percentage of market capitalization. Since we know that firm size is related to returns, using equal-weight rather than value-weight portfolios can skew results (Fama 1998) by including a return-related variable unintentionally. Also, small stocks trade infrequently and market microstructure effects (like bid-ask bounce) can become more of a problem. Lastly, institutional investors are limited in their ability to take a position in smaller stocks due to low dollar trading volume, causing a large market impact for any sizeable trade.

The last major category is related to return reversals. De Bondt and Thaler (1985) find that stocks which performed extremely poorly over the past few years (up to 5) perform relatively better over next few years (up to five) and vice versa. Fama and French (1996) claim that their three-factor model explains this return reversal. Asness (1994) also finds these reversals to be significant. Reversals are also seen on a short scale: Lehmann (1990) and Jegadeesh (1990) find strong negative first-order serial correlation in monthly security returns, leading stocks with strong performance last month to do relatively less well the next month (and vice versa). This effect is very significant (Asness 1994), but it's hard to make a straightforward long-short trading strategy based solely on this because fully turning over a long and short portfolio every month leads to turnover of roughly 2,400% of the total capital per year, leading to very high trading costs.

3. Methodology

We will treat the problem of predicting stock returns as an experts problem, and will show how to incorporate stock features into the above framework. Thus we are looking for performance bounded algorithms. We use two data sources. The first is professor Kenneth French's data library.[8] The second is the center for research in security prices (CRSP). All data are monthly returns, and span the period from 1956-2005 (for data from Kenneth French's data library) 1950 to 2006 (for CRSP data) or 1950-2005 (for industry portfolios created from CRSP data).

3.1. Experts and Online Learning

Our goal will be to find long-only and long-short portfolios that perform well, with positive alpha. As an experts problem, we will look at investment strategies which do almost as well as the best individual asset chosen in hindsight. In different experiments, the assets we use will be different portfolios. In the simple experts framework, there are **N** assets, and we want to invest every period such that our terminal wealth is nearly as good as the wealth achieved by the best asset chosen in hindsight.

Our data consists of a set of returns, $\mathbf{r_j(t)}$, for asset **j** over period **t**. The returns are price relatives with dividends reinvested, so $\mathbf{r_j(t)=1.02}$ means that asset **j** had a return of 2% over period **t**. The cumulative return is the return over the entire horizon for asset **j**

$$R_j(T) = \prod_{t=1}^{T} r_j(t) \tag{3.1}$$

Our goal is to develop an online strategy **S** which picks non-negative weights **w(t)** every period, which sum to one. This strategy's return, over any single period, will be

$$r_S(t) = \sum_{j=1}^{N} \left(w_j(t) \times r_j(t)\right) \tag{3.2}$$

Every period, the strategy can change its weight vector **w**, and its cumulative return will be

$$R_S(T) = \prod_{t=1}^{T} \left(r_S(t)\right) \tag{3.3}$$

The goal of an investment strategy is to minimize log regret. Regret is the factor by which the online strategy's cumulative return lags the cumulative return of the best asset chosen in hindsight. That is, defining

[8] http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

$$R^{OPT}(T) = \underset{j}{MAX}\left(R_j(T)\right) \quad \textbf{(3.4)}$$

we can define a strategy's regret as

$$REGRET = \frac{R^{OPT}(T)}{R_s(T)} \quad \textbf{(3.5)}$$

or its log regret as

$$LOG\_REGRET = \log\left(R^{OPT}(T)\right) - \log\left(R_s(T)\right) \quad \textbf{(3.6)}$$

Machine learning people often think in terms of log regret, since it adds over time and starts from zero. The goal is to develop a strategy with low log regret no matter what sequence of returns occurs. Our performance guarantee will be a regret bound, saying that the regret will never be high. Though it's tempting to try and frame a performance guarantee that promises some high level of return, rather than merely low regret, this is impossible, since the market sequences that occur may not have any opportunities to earn high returns.

3.2. Features and Feature portfolios

We will use the following five stock features: for motivation, see section 2.4 on related work.

1. **PAST(60,13)** – A stock's return over the past 60 months, excluding the last 12 months.
2. **PAST(12,2)** - A stock's return over the past 12 months, excluding the last month.
3. **PAST(1,1)** - A stock's return over the past month.
4. **MKT** – A stock's market capitalization.
5. **BM** – the ratio of a stock's book value to its market value. This measures a stock's "value" properties.

We will combine the experts framework with the features above by creating ***univariate feature portfolios***. This is best illustrated with an example. Suppose we choose to create 10 feature portfolios based on the **BM** feature. To do this, at the beginning of every period we would rank stocks by their **BM** value: those with the highest BM would go in portfolio 1 while those with the lowest **BM** would go in portfolio 10. We do this every period: thus portfolio 1 at any time holds only the 10% of stocks with the highest **BM** value at the beginning of every period. We can just as easily create feature portfolios based on any of these features, or on combinations of these features.

Once we've defined a set of feature portfolios, each feature portfolio becomes an investable asset and we treat it as an expert. Though the actual stocks underlying the feature portfolio will change over time, a feature portfolio still defines a valid investment strategy. Often we'll talk about extreme feature portfolios: these are the feature portfolios with the stocks that most strongly and most weakly exhibit a feature. For example, if we're conditioning on the **PAST(1,1)** feature, the extreme feature portfolios are those which hold stocks with the largest and smallest values of the **PAST(1,1)** feature on any given date (portfolios 1 and 10). That is, the first of these two extreme portfolios would hold stocks which performed extremely well the previous month, and the other would hold stocks which performed extremely poorly the previous month.

4. Universal Portfolio algorithms

A portfolio selection algorithm is said to be "universal" if it achieves the same asymptotic long-term growth rate the best constantly rebalanced portfolio chosen in hindsight (Cover 1991, Blum and Kalai 1997). This is a stronger guarantee than the best individual stock, chosen in hindsight. Borodin, El-Yaniv, and Gogan (2006) point out that any portfolio allocation scheme can be made universal by initially investing some tiny fraction of its wealth in a universal algorithm.[9]

4.1. Constantly Rebalanced Portfolios

A constantly rebalanced portfolio (CRP) is defined by a vector $w \in \Re^{N+}$ whose elements sum to 1. It invests its wealth across N assets according to **w**, and rebalances every period so that, at the beginning of every period, the fraction $\mathbf{w_j}$ of total wealth is invested in each asset **j**. CRPs include individual assets as a special case. The total return of a CRP is

$$R_w(T) = \prod_{t=1}^{T} \left(\sum_{j=1}^{N} w_j \times r_j(t) \right) \qquad \textbf{(4.1)}$$

Though it's not obvious, it's possible for a CRP to have a higher return than any individual stock. From Blum and Kalai (1997), consider the case of two stocks, one of which has returns (1, 1, 1, 1,…) and the other of which has returns (2, ½, 2, ½, …). Holding either individual stock gives a constant return of 1. However, a (½ ½) CRP gives a return of (½ * 2 + ½ * 1) = 3/2 on odd days and (½*½ + ½*1) = ¾ on even days, for a return of 9/8 every two days. Also, CRPs include individual assets as a special case.

4.2. Universal portfolio algorithms

[9] This fact should raise a red flag for the reader that the asymptotic guarantee given by an algorithm's being universal is not a strong guarantee.

A universal algorithm defines its regret with respect to the best CRP chosen in hindsight, rather than just the best asset.
That is, let

$$R_{UNIV}^{OPT}(T) = \underset{w \in S_N}{MAX}\left(R_w(T)\right) \tag{4.2}$$

where S_N is the N-dimensional simplex. Then a universal algorithm defines its regret as

$$REGRET_{UNIV} = \frac{R_{UNIV}^{OPT}(T)}{R_s(T)} \tag{4.3}$$

which is always at least as large as the standard measure of regret. A portfolio allocation algorithm is said to be universal if its "universal" log regret is sub-linear in T for all market sequences. If this is the case, that portfolio will have the same asymptotic rate of return as the best CRP, chosen in hindsight.

The most simple universal portfolio algorithm (Cover 1991, Blum and Kalai 1997) simply invests "smoothly" in all CRPs on the probability simplex, and buys and holds each CRP. I will refer to this as "the Universal Portfolio algorithm" (UNIV) and it achieves

$$REGRET_{UNIV} \leq (T+1)^{N-1} \tag{4.4}$$

Another universal portfolio algorithm is the Online Newton Step (Agarwal, Hazan, Kale, and Schapire 2006). This makes a second-order approximation to the current loss function (that is, loss as a function of location on the simplex) and gives a regret bound of

$$REGRET_{UNIV}(ONS) \leq \exp\left(22 \times N^{1.5}\sqrt{T \times \log(N \times T)}\right) \tag{4.5}$$

Agarwal et. al. use "regret" to refer to what I call "log regret", so the bound in equation 4.5 looks different from the one in their paper.

4.3. Properties of CRPs

CRPs are a little weird. For example, a constantly rebalanced portfolio will continually invest more and more wealth in a falling stock. CRPs have the following properties:

Property 1: The terminal wealth of CRPs is concave in **w.**
Proof: The terminal log wealth in a CRP is

$$\log(R_w(T)) = \log\left(\sum_{t=1}^{T}\left(\sum_{i=1}^{N-1} w_i \times r_i(t) + \left(1 - \sum_{i=1}^{N-1} w_i\right) r_N(t)\right)\right) \tag{4.6}$$

$$\frac{d}{dw_i}(\log(R_w(t)) = \left(\frac{1}{R_w(T)}\right)\left(\sum_{t=1}^{T}\left(r_i(t) - r_N(t)\right)\right) \tag{4.7}$$

$$\frac{d^2}{dw_i^2}(\log(R_w(t)) = -\left(\frac{1}{R_w(T)}\right)^2 \left(\sum_{t=1}^{T}\left(r_i(t) - r_N(t)\right)\right)^2 \leq 0 \qquad \textbf{(4.8)}$$

Because of property one, we can easily find the best CRP in hindsight using a gradient-descent or coordinate descent method.

CRP's also have inherent turnover. This is driven by *cross-sectional* (as opposed to time series) variation in individual stock returns: if there is a lot of variation in a given month for a set of stocks, a CRP investing in those stocks will have to do a lot of buying and selling to maintain its balance. Since turnover incurs transaction costs, this is a disadvantage of CRPs versus buy and hold strategies.

5. Generalized buy-and-hold algorithms

The procedure of initially investing equal-weight in each asset, and then simply holding each asset (reinvesting dividends) is the buy-and-hold algorithm. This is NOT the same as a constantly rebalanced portfolio, since **w** is not constant over time. At any time **t>1**, the investment $\mathbf{w_j(t)}$ in asset **j** for the buy-and-hold algorithm is given by

$$w_j(t) = \frac{R_j(t-1)}{\sum_{k=1}^{N} R_k(t-1)} \qquad \textbf{(5.1)}$$

More importantly, the cumulative return to the buy-and-hold strategy is

$$R_{BH}(t) = \frac{1}{N}\sum_{j=1}^{N} R_j(t) \qquad \textbf{(5.2)}$$

Common sense tells us that this is true, but we can prove it by induction. It's true for t=0 since all cumulative returns are one at time zero. For t>0, we use the inductive hypothesis to show

$$R_{BH}(t) = r_{BH}(t) \times R_{BH}(t-1) = \left(\sum_{j=1}^{N} r_j(t) \times w_j(t)\right) \times \frac{1}{N}\sum_{j=1}^{N} R_j(t-1) \qquad \textbf{(5.3)}$$

and then use the definition of $\mathbf{w_j(t)}$ given above in EQ 5.1 to arrive at equation 5.2. From EQ 5.3 and the non-negativity of returns, we get the regret bound

$$\text{REGRET(BH)} \leq N \qquad \textbf{(5.4)}$$

or

$$R_{BH}(T) \geq \frac{1}{N} R^{OPT}(T) \tag{5.5}$$

This bound is much stronger than the bounds for universal portfolio or for online Newton step, but it's a weaker comparison class: the buy-and-hold algorithm's regret is only compared to the best asset, while Online Newton Step and Universal Portfolio compare to the best CRP. It's an empirical question whether the larger companion class or the tighter bound matters more.

5.1. Geometric Mean Returns

The most commonly used measure of return is the geometric mean return. This is defined as

$$\mathrm{ret}(T) = \left(R(T)\right)^{1/T} \tag{5.6}$$

Using EQ 5.5, we can bound the geometric mean return achieved by the buy-and-hold algorithm as

$$\mathrm{ret}_{BH}(T) \geq \left(\frac{1}{N}\right)^{1/T} \mathrm{ret}^{OPT}(T) \longrightarrow \mathrm{ret}^{OPT}(T) \tag{5.7}$$

where $\mathbf{ret^{OPT}(T)}$ is the geometric mean return of the best asset. Thus the buy-and-hold algorithm performs very nearly as well as the best individual asset, chosen in hindsight.

5.2. Buy and hold optimality

No other online strategy will have as strong of a performance bound as buy-and-hold, when regret is measured against the best asset.

Property 5.1: For $\varepsilon > 0$,no online strategy has worst-case bound

$$R_S(T) > \left(\frac{1}{N} + \varepsilon\right) R^{OPT}(T)$$

Proof: Consider the following sequence of returns. On date 1, a random asset **j** has a return of 1 and all other assets have a return of 0 (or **ε/N**). On all future dates, all assets have a return of 1. The optimal strategy, chosen in hindsight, fully invests in stock **j** and has cumulative return 1. Any strategy which is initially equal-weight has cumulative return of **1/N** (or **1/N + ε**). Any strategy which isn't initially equal-weight at the beginning of period 1 can have cumulative return less than **1/N** (or **1/N + ε**) if it underweights asset **j**. Thus no strategy can guarantee regret less than **N**.

We can actually say something stronger. The proof used in property 5.1 was basically a one-period example.

Property 5.2: Any strategy other than buy-and-hold has worst-case regret > **N**.
Proof: Consider comparing buy-and-hold (BH) with some other online strategy **S**. Let **t'** be the period where S first makes an allocation decision differently from **BH**. Then, at the beginning of period **t'**, both **S** and **BH** have the same amount of total wealth, and S will underweight some stock **j** relative to buy-and-hold. If, over period **t'**, all stocks other than stock j have returns of 0 and stock **j** has a return of 1, buy-and-hold's terminal wealth will be $R_j(t)/N$, as guaranteed by the regret bound. Strategy S will have strictly lower wealth, since it had the same wealth as BH up to time **t'** but then it underweighted the one stock with non-zero return. Thus buy-and-hold is strictly worst-case optimal.

5.3. Generalized Buy and Hold for other portfolio types

We can generalize the buy-and-hold algorithm for different types of portfolios. We'll consider short-portfolios, which are required to only short assets, and long-short portfolios, which are always long and short equal dollar amounts of assets.

5.3.1. Short-only portfolios

Here we will work with portfolios which only short assets. The focus here is theoretical, as such a portfolio probably wouldn't be used in practice. From here we will make the assumption that asset returns are never greater than 2: otherwise short positions can lose an unbounded amount of money. Depending on what assets make up our universe and what periods of time we use, this may be very reasonable or very restrictive. If we define our periods to be months and our assets to be a set of somewhat diversified portfolios, we're probably safe since very few diversified portfolios will double in a month. However, if we define our periods to be years and our assets to be high volatility stocks, we may be in trouble, since a stock can easily double in a year. Also, we'll make the (false) assumption that there are no borrowing costs. In this case, shorting a stock on date **t** means that we receive the date **t** price of the stock, but we must deliver a share of the stock in the future.
Consider the case where we have \$1 to invest, and our broker is willing to give us a leverage ratio of 1. If we short \$1 of stock **i** on date **t**, we ***receive*** a dollar, but to close the position on date **(t+1)** we have to spend $\mathbf{r_i(t)}$ to buy shares to cover our short, so our total wealth on date **(t+1)** consists of (1) our original dollar

(2) the dollar we received from the short (3) the liability of $\mathbf{r_i(t)}$ because we owe a share. Thus, writing short return with $\mathbf{r^s}$, our short return is

$$r_i^s(t) = (2 - r_i(t)) \tag{5.8}$$

Like with long position, we can write the wealth on date **T** of continuously shorting stock **i** using a product, as long as we reinvest all profits every period

$$R_i^s(T) = \prod_{t=1}^{T} r_i^s(t) = \prod_{t=1}^{T} (2 - r_i(t)) \tag{5.9}$$

This equation correctly gives the terminal wealth of continuous shorting, but to achieve it we have to act differently than with long-only positions. Specifically, it says that if our short position makes money, we have to increase our short position and short more shares: if it loses money, we decrease the position. We call this an ***active short position.*** The "optimal" **BH** strategy for short-only portfolios involves these active shorts. It should make sense that equation 5.9 can only be correct if we somehow increase our bet after being successful: if we simply short one share and wait **T** periods to buy it back, we can't possibly do more than double our wealth, since the stock can't go below zero. But equation 5.9 says we can increase our wealth by an arbitrary factor: to do this, we increase our position every time the short makes money. Thus short positions are inherently different from long positions: with long positions the change in stock price ***automatically*** adjusts how our "bet" is distributed across stocks. With short positions, we have to do this adjustment ourselves. This leads to an interesting observation.

Property 5.4 The returns to an ***active*** short position depend on the price path of the stock.

Proof. Consider two stocks initially priced at 1024, neither of which pay dividends. Stock one drops to 1 on the first day (for a return of **.0009766**), and on the next nine days it has a return of 1 (i.e. a constant price). For stock 2, the stock price halves each day, giving it a return of **0.5** every day.
On stock 1, an active short position earns

$$R_1^s(10) = \prod_{t=1}^{10} (2 - r_1(t)) = (2 - .0009766) \times 1 \approx 1.999$$

On stock 2, an active short position earns

$$R_2^s(10) = \prod_{t=1}^{10} (2 - r_2(t)) = (2 - .5)^{10} \approx 57.665$$

The difference is because, with stock 2, every time it drops we use the profit we made to increase our short position.

If we define $R^{s,OPT}(T) = \underset{i}{Max}\left(R_i^s(T)\right)$ as the highest return from a short position in any asset, then bounds exactly the same as EQ 5.5 and EQ 5.7 will hold, and we won't be able to do better than initially shorting every stock in equal proportions and maintaining that in an active short position.

5.3.2. Long-Short portfolio management

We could hope to achieve higher terminal wealth by managing a long and a short portfolio together. Let $\mathbf{R^{LS}(T)}$ be the cumulative return achieved by managing some long-short portfolio. A fixed long-short portfolio is one that puts a fixed weight $\mathbf{w_i}$ on every stock: a positive weight indicates a long position and a negative weight indicates an ***active*** short position. To limit the capital employed, we want

$$\sum_{i=1}^{N} w_i = 0 \quad \text{and} \quad \sum_{i=1}^{N} |w_i| \le 2 \qquad \textbf{(5.10)}$$

When we take a position **w** in a long-short portfolio, we short those stocks with negative **w** and go long those with positive **w**. We can use the standard trick of writing

$$w = w^+ - w^- \qquad \textbf{(5.11)}$$

where $\mathbf{w^+}$ is **MAX(w,0)** and w- is **MAX(-w,0)**. Then we have a combination of a long and a short portfolio: our long portfolio will return

$$\sum_{j=1}^{N} \left(w_j^+ \times r_j(t)\right) \qquad \textbf{(5.12)}$$

and the short portfolio will return

$$\sum_{j=1}^{N} \left(2 - w_j^- \times r_j(t)\right) \qquad \textbf{(5.13)}$$

But there's one caveat: the formulas in equations 5.12 and 5.13 both include the original one dollar of capital that we start with: if we just add these two formulas we'll be double-counting that dollar. So the correct formula for the return of a long-short portfolio adds equations 5.12 and 5.13 and subtracts one, giving

$$r_w^{LS}(t) = \sum_{j=1}^{N}\left(w_j^+ \times r_j(t) + \left(2 - w_j^- \times r_j(t)\right)\right) - 1 = \sum_{j=1}^{N}\left(1 + w \times r_j(t)\right) \quad \textbf{(5.14)}$$

Using this, the date **T** cumulative return of a long-short portfolio is

$$R_w^{LS}(T) = \prod_{t=1}^{T}\left(1 + \sum_{i=1}^{N} w_i r_i(t)\right) \quad \textbf{(5.15)}$$

To compute regret, we need a “best expert” to compare too. With long-short portfolios, we compare to the best long-short pair chosen in hindsight. The date **T** cumulative return of the optimal long-short pair, chosen in hindsight, is

$$R^{LS,OPT}(T) = \underset{i,j}{\text{Max}}\left[\prod_{t=1}^{T}\left(1 + r_i(t) - r_j(t)\right)\right] \quad \textbf{(5.16)}$$

Because a short position is involved, achieving these returns will involve actively rebalancing. Thus in some sense it’s more than a “buy and hold” strategy. Also, achieving the returns in a long-short portfolio involves the following steps:

1) At the beginning of each period, short the assets with a negative weight
2) Use the proceeds of this short sale to purchase the assets with a positive weight
3) Put the original capital in a margin account (this will earn a low interest rate)
4) At the end of the period, readjust all positions to reflect changes in wealth: this is required whether or not the weights are fixed.

We include long-short portfolios for two reasons. First, from a statistical learning point of view, as far as we know no regret bounds for long-short portfolios exist. Second, from a finance point of view, the returns to a long-short portfolio more directly reflect investment skill.

Long-short returns are costless: the proceeds from the short sale are used to finance the long position. Thus they don’t include a risk-free rate, so they don’t include compensation for time-value of money. So if a long-short portfolio earns positive returns, either the assets in the long portfolio are riskier than those in the short portfolio, or the return is alpha (skill-based). We’ve seen long-short portfolios before: In the CAPM, the market portfolio can be thought of a portfolio that is long the market and short the risk-free rate.

One caveat is that throughout the discussion of long-short portfolios, we’ll have to assume that the difference of returns between any two assets is not greater than one. Otherwise long-short portfolios can lose an unbounded amount of wealth, equation 5.15 will not be correct, and no useful bounds will hold.

5.3.3. Theoretical properties of long-short portfolio

Besides being intuitively useful for decomposing returns differently, long-short portfolios can achieve asymptotically higher growth rates than either long-only or short-only portfolios. We can see this when one stock always has a negative return and the other always has a positive return. Here shorting the negative one and going long the positive one will give a growth rate larger than a long-only or short-only portfolio, since the long-short portfolio's return will be related to the spread between these positive and negative returns.

Long-short portfolios are distinct from a separately-managed long portfolio and short portfolio. For one thing, a long-short portfolio is constrained to have the same amount of capital long and short. But more importantly, in a long-short portfolio the profits every period from both the long and short sides are pooled together.

Intuitively, if we have a good way to separate stocks (or groups of stocks) into ones that tend to have low returns and ones that tend to have higher returns, then long-short portfolios will be useful and let us earn the return "spread" with no invested capital.

5.3.4. Low Regret long-short strategy

There is a generalized buy-and-hold (BH) strategy for long-short portfolios, which achieves

$$R^{LS,BH}(T) \geq \left(\frac{1}{N^2}\right) \times R^{LS,OPT}(T) \qquad \textbf{(5.17)}$$

This gives a geometric mean return bound of

$$ret^{LS,BH}(T) \geq \left(\frac{1}{N^2}\right)^{1/T} \times ret^{LS,OPT}(T) \qquad \textbf{(5.18)}$$

thus giving asymptotic growth rate as good as the best pair of assets.

The strategy will invest in ***pairs*** of assets, including self-pairs. That is, for every pair of assets **(i,j)**, the strategy will initially allocate $\mathbf{1/N^2}$ of its capital to going long stock **i** and short stock **j**. Some wealth will be allocated to self-pairs: this is OK since that wealth is essentially held out of the market. After this initial allocation, it will let the wealth in each pair compound at that pair's growth rate: this is essentially buying and holding each pair. However, because the shorts are active shorts, some transactions will need to be made. At any time **t > 1**, the weight on pair **(i, j)** will be

$$w_{i-j}(t) = \frac{\prod_{\tau=1}^{t-1}\left(1 + r_i(\tau) - r_j(\tau)\right)}{\sum_{a,b}\left(\prod_{\tau=1}^{t-1}\left(1 + r_a(\tau) - r_b(\tau)\right)\right)} \quad \textbf{(5.19)}$$

However, we can more easily do analysis by noting that the wealth in each pair, per initial dollar of capital, is

$$\$_{i-j}(t) = \frac{1}{N^2}\prod_{\tau=1}^{t}\left(1 + r_i(t) - r_j(t)\right) \quad \textbf{(5.20)}$$

thus achieving the bound 5.17 by the non-negativity of returns.

Recall the long-short portfolio requirements: $\sum_{i=1}^{N} w_i = 0$ and $\sum_{i=1}^{N} |w_i| \le 2$.
These requirements relate to the holdings of assets, not the holdings of pairs. We automatically meet these requirements by treating each pair of assets as an asset itself, since each pair is automatically long and short equal dollar amounts. And since the sum of the pair weights $\mathbf{w_{i\text{-}j}}$ (which are always positive) is never more than one, the sum of the absolute values of the asset weights $\mathbf{w_j}$ is never more than two.

By treating each ***pair*** of assets as a single asset, we never have to deal directly with the long and short holdings of each asset. This is good, because these individual asset holdings do not lie on a simplex: they lie on a strange shifted sum of simplexes. However, the portfolio weights on each ***pair*** of assets do lie on a simplex. Interestingly, this low-regret strategy initially makes zero net investment, since at first the holdings in all pairs cancel out (i.e. a position in pair "long asset **i** short asset **j**" cancels a position in pair "long asset **j** short asset **i**". Over time, as wealth in each pair grows or shrinks, the algorithm will take an active position.

Property 5.3.4.1: The bounds in equations 5.17 and 5.18 are tight, in the sense that there exists a sequence of returns where the pair strategy achieves only slightly above this bound.

Proof: Consider a sequence of asset returns where asset 1 always returns 1.02, asset 2 always returns 0.98, and assets 3 to N always return 1. Here the pair returns are as follows, where **X** is any asset from **3** to **N**:
$r_{1-2}(t) = 1.04$, $r_{1-X}(t) = 1.02$ $r_{2-1}(t) = 0.96$ $r_{2-X}(t) = 0.98$, $r_{X-1}(t) = 0.98$, and $r_{X-2}(t) = 1.02$. Of course all self-pairs have returns of 1.
Thus there are:
$\mathbf{(N\text{-}2)^2 + 2 = N^2 - 4N + 6}$ pairs with a return of **1**

2(N-2) pairs with a return of **1.02**
2(N-2) pairs with a return of **0.98**
1 pair with a return of **0.96**
1 pair with a return of **1.04**

Since wealth is simply allocated to each pair initially and allowed to grow, the total wealth at time T is:

$$R^{LS,BH}(T) = \frac{1}{N^2}\left(1.04^T + 2(N-2)1.02^T + 2(N-2)0.98^T + 0.96^T + N^2 - 4N + 6\right)$$

Thus for large T, $\lim_{T\to\infty}\left(R_S^{LS}(T)\right) = \frac{1.04^T}{N^2} = \frac{1}{N^2}\left(R^{LS,OPT}(T)\right)$

We are not sure if a strategy exists to achieve a significantly tighter bound than equation 5.17: this is an open question. We know it's possible to replace the $\mathbf{N^2}$ with **N(N-1)+1** in equations 5.17 and 5.18 by allocating $\frac{1}{N(N-1)+1}$ to each pair of stocks that isn't a self pair and holding the same amount out of the market: this works because all **N** of the self-pairs all have a return of 1. We can also weaken the comparison class to exclude self pairs, defining the best pair as

$$R^{LS,OPT}(T) = \underset{i \neq j}{\mathrm{Max}}\left[\prod_{t=1}^{T}\left(1 + r_i(t) - r_j(t)\right)\right] \qquad \textbf{(5.21)}$$

This is what we do in our experiments, and this allows the bound to be **N(N-1).**

5.3.5. Hidden issues in shorting

Until now we've assumed that our shorts, or the short portions of our long-short portfolios, never have negative wealth. Here we show why.

Property 5.3.5.1: When returns are unbounded for long-short portfolios and **T** can be arbitrarily large, no strategy which doesn't use foresight can achieve a bound

$$R^{LS,S}(T) \geq g(N)R^{LS,OPT}(T) \qquad \textbf{(5.22)}$$

where **g(N)** is any strictly positive finite function of **N**.

Proof:
Consider a series of asset returns where, at first, asset 1 always returns **1.02,** asset 2 always returns **0.98**, and assets 3 to N always return **1**. Let **S** be any

non-anticipating strategy. Note that the best long-short pair has cumulative return $R^{LS,OPT}(t) = 1.04^{t}$. Consider possibilities for **S**.

1) **S** never has **w2 < -0.5**. In this case, **S**'s rate of return is never greater than 3%, and as T becomes large **S** will lag the best-pair's cumulative return by an exponentially increasing amount in **T**. Formally,

$$R^{LS,S}(T) \leq 1.03^{T} \Rightarrow \quad \textbf{(5.23)}$$

$$\frac{R^{LS,S}(T)}{R^{LS,OPT}(T)} \leq \left(\frac{1.03}{1.04}\right)^{T} \leq g(N) \text{ (for some } \mathbf{T}\text{)} \quad \textbf{(5.24)}$$

2) At some time **t'**, **S** sets its allocation so that **w2 < -0.5**. In this case, if asset 2's return the next period is **4**, **S** will have a tremendous debt from its short position in asset 2 and will end with negative cumulative return.

A similar result holds for pure short portfolios.

5.4. Hedge

The Hedge algorithm (Freund and Shapire 1995) develops a weight update scheme to solve the experts problem. It is another performance bounded algorithm. In Hedge, a number $\beta \in (0,1)$ is chosen and used in the weight update rule. Every period, each of the **N** experts suffers some loss $\ell(t) \in [0,1]$. Hedge chooses a weight allocation scheme where, initially, the algorithm bets equal weight on every expert. After every period, the weight on each expert is multiplied by a factor $w_i(t+1) = w_i(t) \times \beta^{\ell_i(t)}$, and then all weights are normalized to sum to one. Then, where $\mathbf{L_{HEDGE}}$ is the cumulative loss of the hedge algorithm, this assures that

$$L_{HEDGE} \leq \frac{L^{OPT} \times \ln\left(1/\beta\right) + \ln(N)}{1-\beta}$$

where $\mathbf{L^{OPT}}$ is the loss of the expert with minimal loss.

Property 5.4.1: Using $\ell_i(t) = -\ln(r_i(t))$ and $\beta = e^{-1}$ the weight allocation chosen by hedge is the same as buy-and-hold.
Proof: Initially all portfolios will be equal-weight. The weight update rule will apply the following:

$$\frac{w_i(t+1)}{w_i(t)} = \beta^{\ell_i(t)} = \left(\frac{1}{\beta}\right)^{\ln(r_i(t))} = e^{\ln\left(\frac{1}{\beta}\right)\ln(r_i(t))} = r_i(t)^{\ln\left(\frac{1}{\beta}\right)} = r_i(t)^{\ln(e)} = r_i(t)$$

With larger β, Hedge will more slowly allocate wealth to portfolios that have performed well. With smaller beta, Hedge will more quickly allocate wealth to portfolios that have performed well.

The derivation of the loss bounds for the hedge algorithm depends on a bounded loss function, and thus does not apply to the investment case (where the loss function $\ell_i(t) = -\ln(r_i(t))$ is unbounded). In any case, the loss bounds for Hedge are far weaker than those we derived for either the long-only or the long-short generalized buy-and-hold strategies.

6. Empirical Results

We first test our various techniques on two "traditional" data sets. The first is a set of 30 industry portfolios, each of which we treat as an individual asset. The data contains monthly returns from 1950 to 2005. The second data set is the 171 stocks with complete return history from 1960 to 2005. The second data set is biased, because an investor in 1960 would not know ahead of time which stocks would survive 46 years later. The industry data set is largely unbiased, since investors do know what industry a company belongs to.

6.1. Cover's Universal Portfolio algorithm

Cover's algorithm is intractable for a large number of assets, so we cannot implement it on the complete set of assets (industry portfolios or stocks). However, we can implement it on all pairs and/or all triplets of assets. We can also compute the cumulative return of the best CRP, chosen in hindsight across all assets, and see what bound Cover's algorithm would guarantees us if we could run Universal Portfolio on all the assets at once.

6.1.1. Universal portfolio on industry data

In Table 1 we run the universal portfolio algorithm on all pairs and triplets of industries, and then average results across pairs (triplets).[10] Averaged over all pairs, we see that the optimal CRP is only marginally better than the optimal industry in each pair or triplet. The universal portfolio algorithm out-performs the equal-weight CRP by just under 0.01% per month. Adding a 0.5% one-way transaction cost reduces this benefit, but the universal portfolio algorithm still comes out marginally ahead. The optimal industry (from each pair, averaged over pairs) did nearly as well as the optimal CRP. The rightmost column shows the minimum return guaranteed by the lower bound on the performance of the universal portfolio relative to the optimal CRP: this bound is lower than the risk-free rate, but isn't ridiculously low. The results for universal portfolios over

[10] Of course we want to run it on the whole set of industries at once, but we can't because the algorithm is intractable.

triplets of industries are roughly the same: the average universal portfolio beats an equal-weight CRP by a miniscule 0.013% per month, before transaction costs. Because a higher number of assets weakens the bound, the guaranteed return for triplets of industries is actually negative.

Remember that the difference between Cover's universal portfolio algorithm and our generalized buy-and-hold algorithm is that universal portfolio has a weak regret bound with respect to the best CRP, and we have a strong regret bound with respect to the best individual asset. Thus the only economic reason to prefer a universal portfolio algorithm is if the best CRP is predicted to be significantly better than the best individual asset.

Table 2 shows the return of the best CRP, across all industry portfolios, and the best individual industry. The difference is negligible. It also shows the return bound for the universal portfolio algorithm (in the Best CRP column) if it was run on all industry portfolios together (we can't do this since it would take exponential time). The table also shows the return bound for the buy and hold algorithm. Given the performance of the best CRP, the universal portfolio algorithm is sure to lose no more than 23.6% PER MONTH, while the buy-and-hold algorithm will earn at least 0.65% per month. Clearly the universal portfolio bound is orders of magnitude short of being useful.

Universal Portfolio results for all pairs of industry portfolios

Geometric Mean Returns	Universal	Best CRP	Best Industry	Equal Weight	Guaranteed
No Transaction Cost	0.967%	1.011%	1.004%	0.958%	0.04%
0.5% Transaction Cost	0.961%	1.007%	1.004%	0.958%	0.03%
Wealth					
No Transaction Cost	735.95	990.38	941.99	683.14	1.47
0.5% Transaction Cost	707.1	965.35	941.99	683.14	1.41

Universal Portfolio results for all triplets of industry portfolios

Geometric Mean Returns	Universal	Best CRP	Best Industry	Equal Weight	Guaranteed
No Transaction Cost	0.978%	1.045%	1.034%	0.965%	-0.79%
0.5% Transaction Cost	0.970%	1.040%	1.034%	0.965%	-0.799%
Wealth					
No Transaction Cost	759.16	1203.86	1102.2	683.14	0.01
0.5% Transaction Cost	721.06	1156.65	1102.2	683.14	0.01

Table 1

For all pairs / triplets of industries, we compute the return of Cover's universal portfolio algorithm (column 1), the return of the best CRP chosen in hindsight (column 2), the best individual industry for that pair (column 3), the return of an equal-weight portfolio (column 4) and the lower bound on the return of the universal portfolio, which is related to the return of the best CRP. The table reports these results, averaged over all pairs / triplets.

Best CRP using ALL industries

Long Portfolios	Best CRP	Best Industry	Buy and Hold
Geom. Mean Returns	1.18%	1.14%	0.98%
Cumulative Return	2711.3	2056.7	683
Guaranteed Cumulative Return	2.6348E-79		84.47
Guaranteed Geom. Return	-23.60%		0.65%

Table 2

For all 30 industry portfolios, we compute the best CRP chosen in hindsight and the best industry. We also report the performance of the buy-and-hold algorithm. The third and fourth row give performance bounds for Cover's universal portfolio algorithm (under the column Best CRP) and the buy-and-hold strategy.

6.1.2. Universal Portfolio on individual stocks

We also tried Cover's universal portfolio algorithm on returns for individual stocks. Here we used all 171 stocks with a complete return history from 1960 to 2005. Again, we were only able to run on pairs of stocks, since the algorithm is intractable. Looking at Table 3, the universal portfolio algorithm has a slightly larger advantage: it beats the equal-weight portfolios by 0.04%, on average, over pairs of stocks. But the bounds are still weak. Table 4 computes the best CRP over all 171 stocks, and looks at the best individual stock. It also shows the performance bounds for the universal portfolio algorithm run over all stocks at once, and the performance bounds for the buy-and-hold algorithm. The guarantee for the universal algorithm is astronomically weak: it will return no less than -85% per month. The buy-and-hold algorithm's guarantee is 0.58% per month, or over 7% per year. This is very strong, considering the large number (171) of assets under consideration.

All Pairs	**Universal**	**Bets CRP**	**Best Stock**	**Equal-Weight**	**Universal Bounds**
Geometric Returns	0.94%	1.02%	0.98%	0.90%	-0.13%
Cumulative Returns	277.84	465.4	405.63	245.86	0.84

Table 3

For all pairs of stocks, we compute the return to the universal portfolio algorithm (col 1), the best CRP for that pair (col 2), the best stock for that pair (col 3), the equal-weight return (col 4) and the return guarantee from the universal portfolio performance bounds (col 5). The table reports these results, averaged over all pairs.

All 171 Stocks	Best CRP	Best Stock	Buy and Hold
Geometric Returns	1.57%	1.52%	1.00%
Cumulative Returns	5453.73	4148.6	245.85
Guaranteed Geom. Return	-85%		0.58%
Guaranteed Cum. Return	0		24.26

Table 4

Reports results for the best CRP, best individual stock, and buy and hold portfolio for all 171 stocks with complete return history. All returns are monthly. The third column shows the return guarantees for the universal portfolio algorithm (col 1) and for the buy-and-hold strategy (col. 3).

6.1.3. Universal portfolio conclusions

This section's results make a strong case that the bounds given by the universal portfolio technique are economically meaningless, while those given by the buy-and-hold strategy are quite strong (though one of the data sets had selection bias). The reason is that universal portfolios measure their regret versus a much "harder" class of experts (the best CRP), which gives a much weaker bound. In our data sets, the best CRP did not do significantly better than the best individual asset, so the two measures of regret were similar. But even if we had found CRPs in our data set with very good performance, the universal portfolio bounds still would have been weak. The take-home point is that ***asymptotic bounds may not be useful***, and as investors with finite lives we might need something stronger.

The universal portfolio techniques aren't bad: they outperformed the equal-weight strategy by a tiny amount on pairs of assets. But, fundamentally, the only reason to use a performance bounded strategy is if it gives meaningful bounds. If it doesn't, it is simply an arbitrary way to reallocate wealth with risk and return characteristics that aren't well understood, which may include unintentional factor bets.[11] Without meaningful bounds, we might as well use a heuristic strategy instead of a performance bounded strategy.

6.2. Online Newton Step

Online Newton Step (Agarwal et. al. 2006) is another "universal" algorithm, with regret sub-exponential in **T**. However, it is tractable and can be run on large numbers of assets.

6.2.1. Online Newton Step on Industry and Stock Data

[11] Cover's universal portfolio algorithm does well when returns have negative autocorrelation. You'll notice that in the toy example presented in section 4.1 that one of the assets has strong reversals.

Despite being able to run with a far larger number of stocks, online Newton step gives results similar to Universal Portfolio on the industry dataset. As shown in Table 5, it performs a tiny bit better than buy and hold on the set of all thirty industry portfolios, and actually does a little worse when limited to the industry portfolios with the highest or lowest variance. This is surprising since the results presented in the paper by Agarwal et. al. (2006) showed that it performed better with higher-variance stocks.

Looking at Table 6, we see that online Newton step has significantly better returns when dealing with stocks: it beats the buy-and-hold strategy by 0.18% per month, which is significant. Judging by returns alone, Online Newton Step seems to be economically viable. Also, it performs about 20 basis points above buy and hold (that's about 2.4% per year), before transaction costs, on low-variance stocks. However, a look at the standard deviation column shows that Online Newton Step is far riskier, and on a risk-adjusted basis the small increment in return (before transaction cost) that it gives over buy and hold isn't worth the extra variance. Still, it's quite surprising that its returns are so good.

Despite the strong returns on stock data, we see from the bounds in Table 5 and Table 6 that the return guarantees given by Online Newton Step are economically meaningless. We can't even store the guaranteed terminal wealth on a computer (unless we use a big number package or log-based arithmetic) because it under-flows. Also, while Online Newton Step has a big advantage over buy-and-hold (on stock, not industry data), it has a smaller advantage over the equal-weight portfolio. Equal weight portfolios beat buy-and-hold portfolios in data sets with return reversals, and the stocks dataset does exhibit these reversals, so it's likely that Online Newton Step is taking advantage of this.

Online Newton Step on 30 industry portfolios, 1950-2005
Using 672 MONTHLY returns

All Industries (30)	Return	Standard. Deviation	Terminal Wealth
ONS	1.0113%	4.3500%	864.1
Buy and Hold	0.9760%	4.4300%	683.14
Equal Weight	1.0035%	4.3600%	820.6
Best CRP	1.1833%	5.9100%	2711.3
ONS Bound	-100.0000%	NA	10^(-54000)
High Variance Industries (7)			
ONS	0.9910%	5.6200%	755.2
Buy and Hold	0.9685%	5.4400%	650.12
Equal Weight	1.0133%	5.2900%	875.8
Best CRP	1.1391%	6.7900%	2021.5
ONS Bound	-98.9866%	NA	10^(-8197)
Low Variance Industries (8)			
ONS	0.9861%	3.5400%	730.6
Buy and Hold	0.9915%	3.8200%	757.4
Equal Weight	1.0079%	3.6700%	844.9
Best CRP	1.1035%	4.2800%	1595.5
ONS Bound	-99.9222%	NA	10^(-9763)

ONS = Online Newton Step

High variance industries chosen in hindsight with over 1.2 times avg. variance
Low variance industries chosen in hindsight with under 0.8 times avg. variance

Table 5

Reports results for the online Newton step algorithm run on industry portfolios. The fourth row in each set, ONS bound, is the amount of return we are guaranteed by the lower-bound on Online Newton Step's performance.

Stocks (171)	Return	Standard. Deviation	Terminal Wealth
ONS	1.25%	9.00%	943.2
Buy and Hold	1.00%	4.45%	245.85
Equal Weight	1.07%	4.12%	363.1
Best CRP	1.57%	8.41%	5453.8
ONS Bound	-100.0000%	NA	10^(-469808)
High Variance Stocks (33)			
ONS	1.17%	9.25%	604.6
Buy and Hold	1.05%	6.56%	323.68
Equal Weight	1.14%	5.93%	513.4
Best CRP	1.57%	8.44%	5515
ONS Bound	-100.0000%	NA	10^(-55607)
Low Variance Stocks (40)			
ONS	1.17%	4.84%	627.5
Buy and Hold	0.88%	3.48%	128.65
Equal Weight	0.94%	3.62%	179.3
Best CRP	1.11%	4.14%	439.4
ONB Bound	-100.0000%	NA	10^(-71414)

ONS = Online Newton Step
BH = Buy and Hold

Data includes all stocks with a complete return history from 1960 through 2005.
This sample is chosen using hind sight
High variance stocks chosen in hindsight with over 1.4 times avg. stock variance
Low variance stocks chosen in hindsight with under 0.6 times avg. stock variance

Table 6

Reports results for the online Newton step algorithm run on individual stocks. The fourth row in each set, ONS bound, is the amount of return we are guaranteed by the lower-bound on Online Newton Step's performance.

6.2.2. Online Newton Step Conclusions

Though Online Newton Step does have a non-trivial return advantage over buy-and-hold on stock data (not on industry data), it still does not give anything close to economically meaningful bounds. Thus, though it may be useful as a heuristic strategy, it isn't useful as a performance bounded strategy.

6.3. Generalized Buy-and-hold Strategies

We first use the generalized buy-and-hold (BH) strategies on the two sets of assets considered above: the 30 industry portfolios and the 171 individual stocks. Results are presented in Table 7. The long buy-and-hold strategies earn about 1% a month: roughly the market's return. For both the stocks and industries, the bounds are reasonable, but not very tight. The short-only portfolios do uniformly poorly, and the long-short portfolios have negative returns on both datasets. This occurs despite the fact that, in both cases, there is a relatively good pair to go long-short.

All industries

Periods		672	
Industries		30	
		Return	**Wealth**
Long	BH	0.98%	683.1
	Best	1.14%	2056.7
	Bound	0.63%	68.6
Short	BH	-1.19%	0.00033
	Best	-0.97%	0.00140
	Bound	-1.47%	0.00
Long-short	BH	-0.03%	0.8088
	Best	0.37%	11.9433
	Bound	-0.64%	0.01

Stocks with full return data from 1960 to Dec 2005

Periods		552	
Stocks		171	
		Return	**Wealth**
Long	BH	1.00%	245.855
	Best	1.52%	4148.59
	Bound	0.58%	24.26
Short	BH	-1.28%	0.00080
	Best	-0.92%	0.00613
	Bound	-2.20%	0.00000
Long-short	BH	-0.13%	0.4933
	Best	0.77%	69.64
	Bound	-1.97%	0.000017

Table 7

Thus the only time the BH strategies don't lose money is the long-only case, and there they simply earn (roughly) the market return. And we showed earlier that the BH strategies are optimal (or not far from it in the long-short case) in terms of minimizing worst-case regret. This is somewhat disappointing: it would be nice if at least the long-short portfolios could have generated positive returns, since the opportunity is there. But we shouldn't throw in the towel just yet. The bounds in Table 7 are many orders of magnitude tighter than the bounds for the universal algorithms. The reason buy and hold failed is that the returns of individual stocks or individual industries aren't consistent over time. That is, individual stocks don't consistently perform well over long periods.[12] Also, we considered too many experts at once: the bounds will be tighter if we can create a smaller universe of assets.

[12] By this, I mean that we're unlikely to find a stock that beats the market by, say, 1% every month for many decades.

We know that the BH strategies, used on a group of portfolios, will perform reasonably well with respect to the best portfolio. Thus we need to divide up the universe of stocks into portfolios in such a way that some portfolios do very well, and other do very poorly, somewhat consistently. To do this, we'll use feature portfolios.

6.3.1. Feature Portfolios

We start by considering univariate decile feature portfolios, using the features described in section 3.2. Specifically, for each individual feature, we'll divide the universe of stocks into 10 portfolios every period[13]. Portfolio 1 will hold the stocks which most strongly exhibit that feature, and portfolio 10 will hold the stocks which most weakly exhibit that feature. Every portfolio will hold roughly the same number of stocks.

We'll imagine an investor who, at the beginning of the sample period (1955) chooses a feature and believes that feature is related to future stock returns in a monotone way. However, this investor doesn't know what the relationship is. Thus he will use a low regret strategy with respect to the extreme feature portfolios (portfolios 1 and 10) to try to learn this relationship. Annualized returns to these extreme portfolios are shown in Figure 1 through Figure 5. Here we see that, in all cases, the buy and hold (minimum regret) strategy earns very high returns (14% to about 24% per year) and, more importantly, the bound is very tight and economically meaningful! That is, the returns guaranteed by the bound are, in fact, significant, as long as the transaction costs don't eat up the gains.

The results in Figure 6 show the returns achieved by an investor who simultaneously bets on all 10 extreme feature portfolios, using the buy-and-hold strategy to allocate among them. This investor's realized and guaranteed annual returns of around 20% are significant, even though there are ten experts among which to allocate.

We also consider long-short strategies on these feature portfolios. Again, we consider an investor who believes that each individual feature is related to future return in a monotone way. However, this investor will use the long-short generalized buy-and-hold strategy to allocate among extreme feature portfolios. Thus he will allocate between pairs "long portfolio 1 short portfolio 10" and "long portfolio 10 short portfolio 1". Returns for these portfolios are shown from Figure 7 to Figure 11. The absolute magnitude of returns is lower, as it should be since long-short portfolios are costless, so they don't earn a risk-free rate and don't benefit from the general upward market trend. Despite this, the generalized BH strategy earns nearly 10% per year in most cases, and the performance bounds are still quite tight. Figure 12 considers a long-short investor whose strategy can include all 10 extreme feature portfolios (the top and bottom portfolios for each of the five features). This leads to 90 possible long-short pairs, but the investor's

[13] The data for this section are monthly returns from 1955 to 2005 from Professor Kenneth French's data library: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

return and return bound are both near 15% per year despite the large number of experts. This isn't bad for a costless strategy that requires little foreknowledge!

Also, the high returns to the feature portfolios don't seem to compensate for any risk. The charts from Figure 14 to Figure 17 show the beta values for the long-short feature portfolios, and the buy-and-hold strategy, for each of the 5 possible conditioning features. Except for the short-term reversal portfolio, the betas are insignificant (remember, a long-short portfolio with a beta of one should have a return like the market minus the risk-free rate, or roughly 6%). In a few cases (such as feature portfolios conditioning on book to market or PAST(12,2)) the betas of the higher-return long-short portfolios are negative, indicating that not only would they add high returns but they would also be fantastic diversifiers.

The results for generalized buy-and-hold strategies for feature portfolios stand in stark contrast to the results for individual stocks or industries. Feature portfolios generate a set of experts with consistent performance, leading to high returns and very tight performance bounds for online strategies. Also, the generalized buy-and-hold bounds are in all cases economically meaningful, unlike the bounds for the universal portfolio algorithms.

7. Conclusions

We've shown that a simple investment strategy has optimal worst-case regret properties with respect to the best individual asset chosen in hindsight. We showed how to generalize this strategy to long-short portfolio management and still achieve low regret, which (we think) is a new result. We argued that the bounds for "universal portfolio" algorithms, which have a slightly more general notion of regret, are asymptotically strong but prove useless in practice. Finally, we showed how to incorporate features into the experts framework, and how this leads to better realized and guaranteed performance.

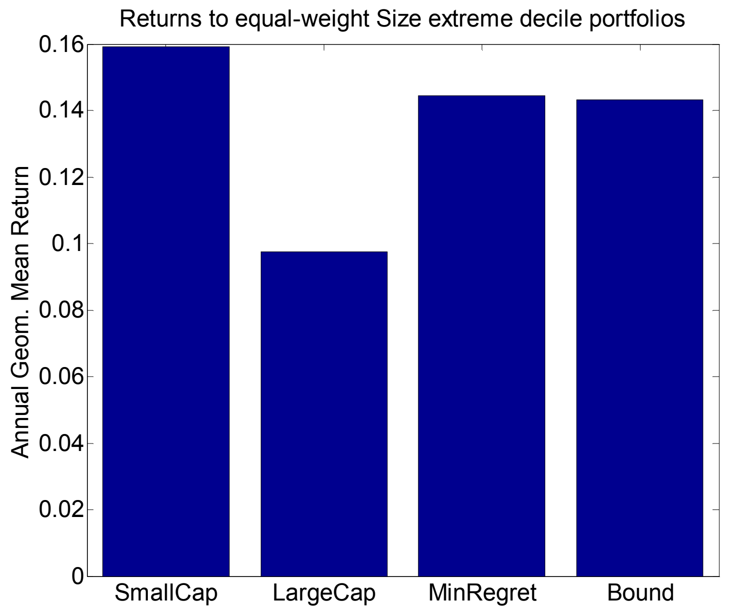

Figure 1

Annual geometric returns to extreme feature portfolio conditioned on market size. The SmallCap portfolio holds only the smallest firms. MinRegret is the return achieved by the buy-and-hold strategy, and Bound is the return guaranteed to the buy-and-hold strategy.

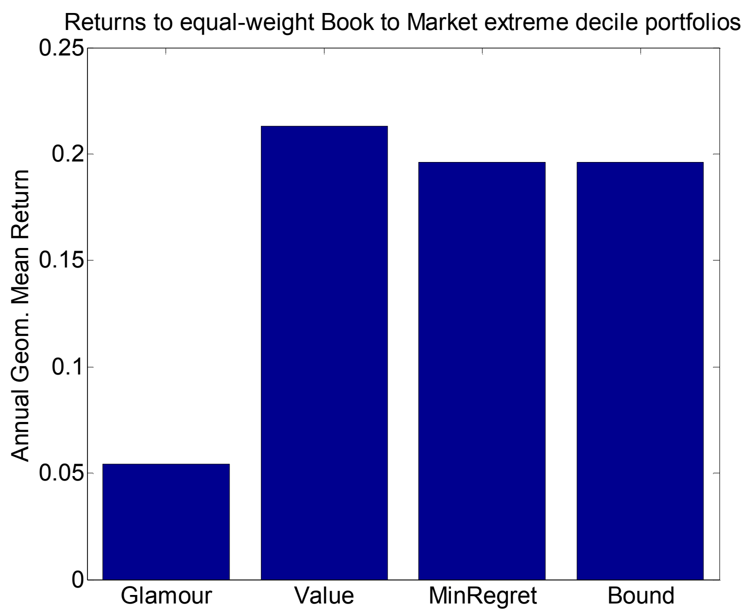

Figure 2

Annual geometric returns to extreme feature portfolio conditioned on book-to-market. The value portfolio holds firms with the highest book-to-market, while the glamour portfolio holds those with the lowest book-to-market. MinRegret is the return achieved by the buy-and-hold strategy, and Bound is the return guaranteed to the buy-and-hold strategy.

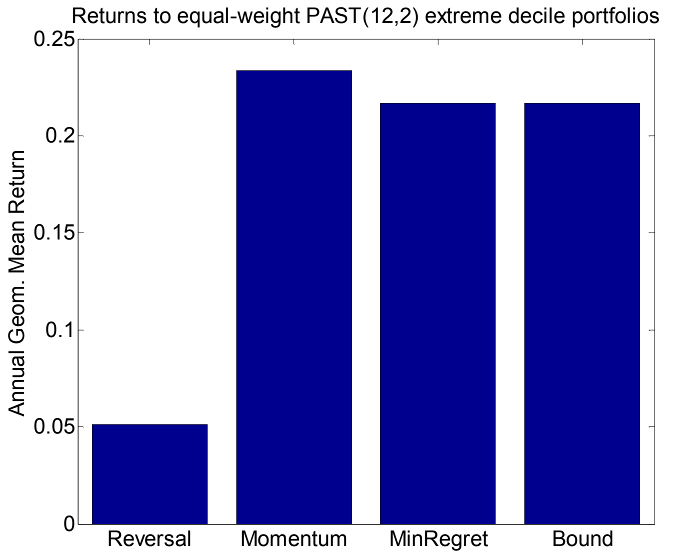

Figure 3

Annual geometric returns to extreme feature portfolio conditioned on PAST(12,2). The Reversal portfolio holds firms with low values of PAST(12,2), and the Momentum portfolio holds firms with high values of PAST(12,2). MinRegret is the return achieved by the buy-and-hold strategy, and Bound is the return guaranteed to the buy-and-hold strategy.

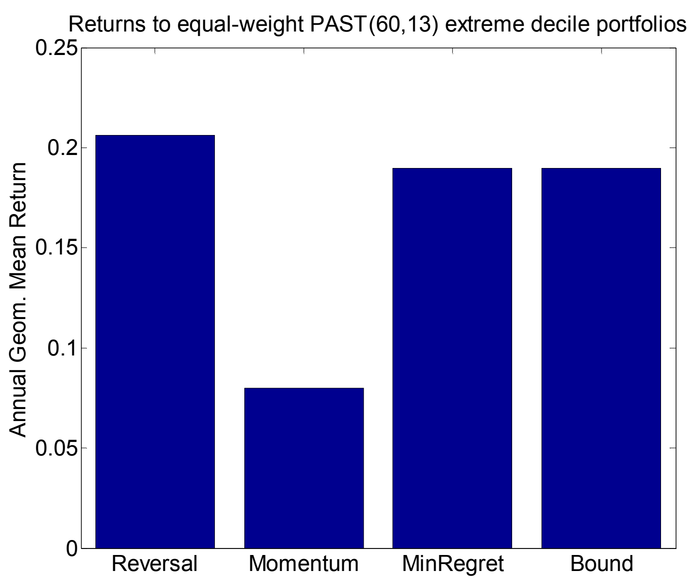

Figure 4

Annual geometric returns to extreme feature portfolio conditioned on PAST(60,13). The Reversal portfolio holds firms with low values of PAST(60,13), and the Momentum portfolio holds firms with high values of PAST(60,13). MinRegret is the return achieved by the buy-and-hold strategy, and Bound is the return guaranteed to the buy-and-hold strategy.

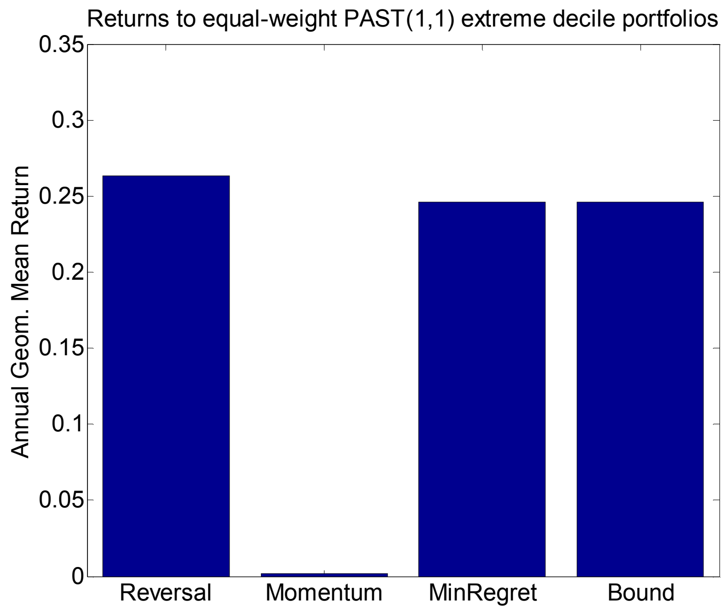

Figure 5

Annual geometric returns to extreme feature portfolio conditioned on PAST(1,1). The Reversal portfolio holds firms with low values of PAST(1,1), and the Momentum portfolio holds firms with high values of PAST(1,1). MinRegret is the return achieved by the buy-and-hold strategy, and Bound is the return guaranteed to the buy-and-hold strategy.

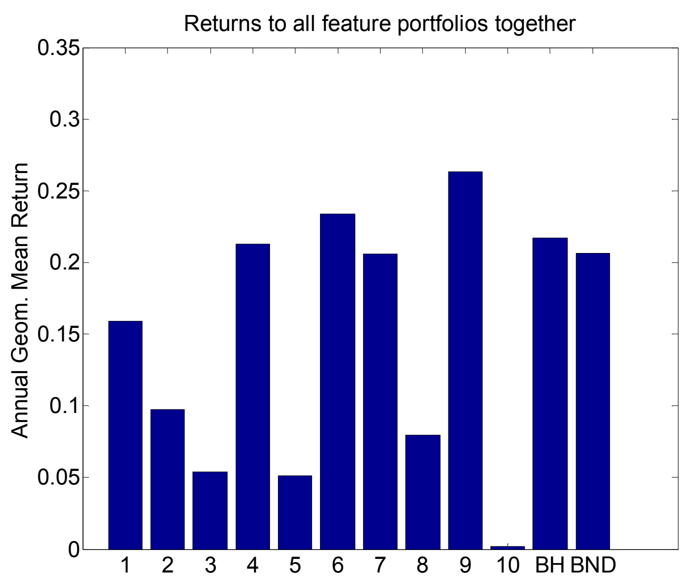

Figure 6

Annual geometric returns to the 10 extreme feature portfolios, two for each of the five features. BH is the generalized buy-and-hold strategy, and BND is the return guaranteed to the buy-and-hold strategy.

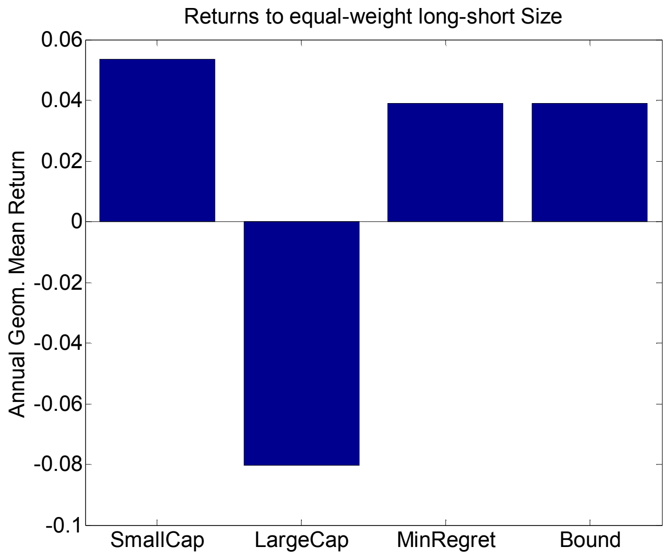

Figure 7

Annual geometric returns to extreme long-short feature portfolios conditioned on Market Cap. The SmallCap portfolio is long the smallest firms and short the largest ones. The LargeCap portfolio is long the largest and short the smallest. MinRegret is the BH strategy, and Bound is the return guaranteed to the BH strategy.

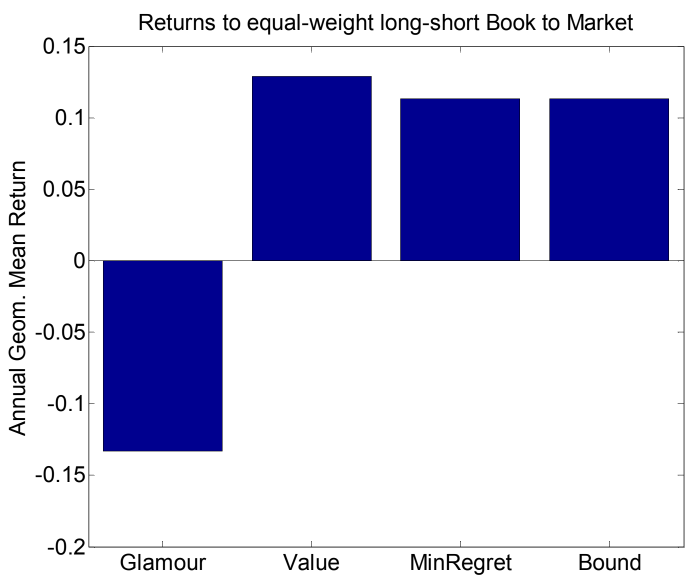

Figure 8

Annual geometric returns to extreme long-short feature portfolios conditioned on Book to Market. The Glamour portfolio is long stocks with the lowest book-to-market and short those with the highest book-to-market values. The value portfolio takes the opposite position. MinRegret is the BH strategy, and Bound is the return guaranteed to the BH strategy.

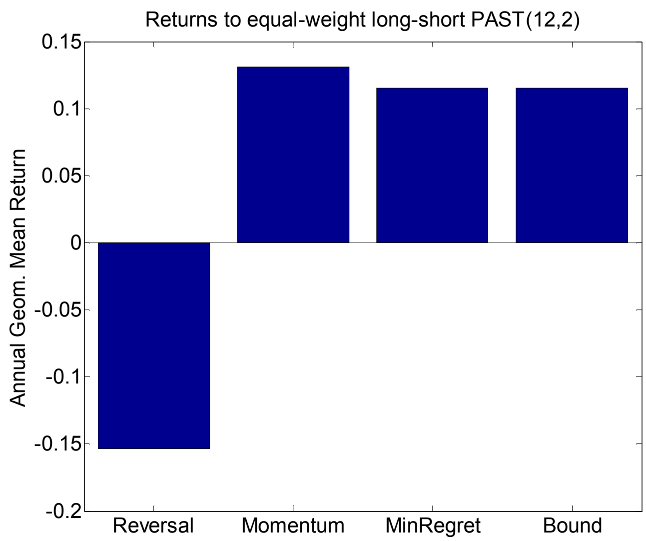

Figure 9

Annual geometric returns to extreme long-short feature portfolios conditioned on PAST(12,2). The Reversal portfolio is long stocks low PAST(12,2) and short stocks with high PAST(12,2). The Momentum portfolio takes the opposite position. MinRegret is the BH strategy, and Bound is the return guaranteed to the BH strategy.

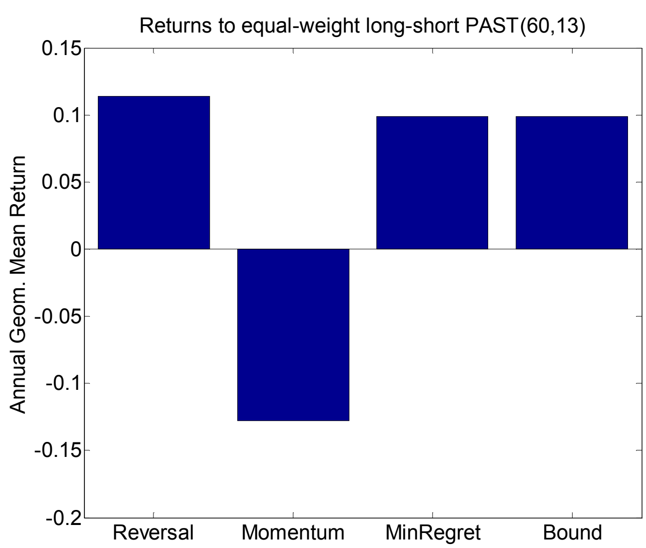

Figure 10

Annual geometric returns to extreme long-short feature portfolios conditioned on PAST(60,13). The Reversal portfolio is long stocks with the lowest PAST(60,13) and short those with the highest PAST(60,13). The Momentum portfolio takes the opposite position. MinRegret is the BH strategy, and Bound is the return guaranteed to the BH strategy.

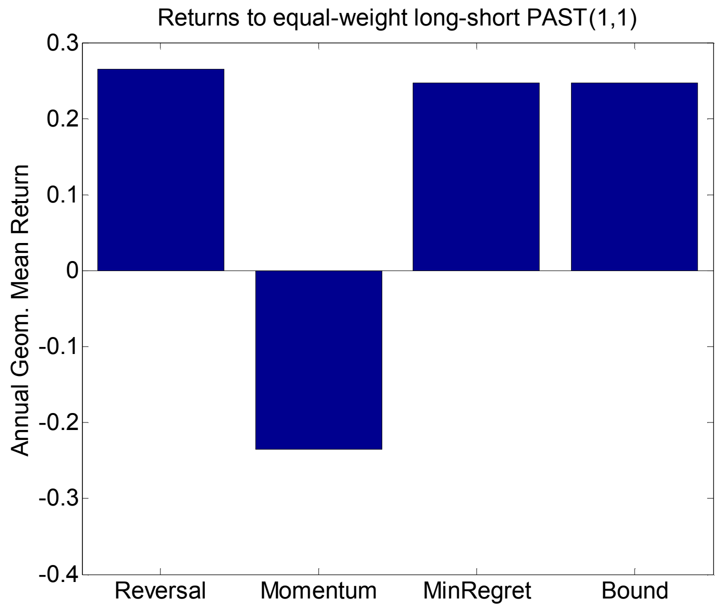

Figure 11

Annual geometric returns to extreme long-short feature portfolios conditioned on PAST(1,1). The Reversal portfolio is long stocks with the lowest PAST(1,1) and short those with the highest PAST(1,1). The Momentum portfolio takes the opposite position. MinRegret is the BH strategy, and Bound is the return guaranteed to the BH strategy.

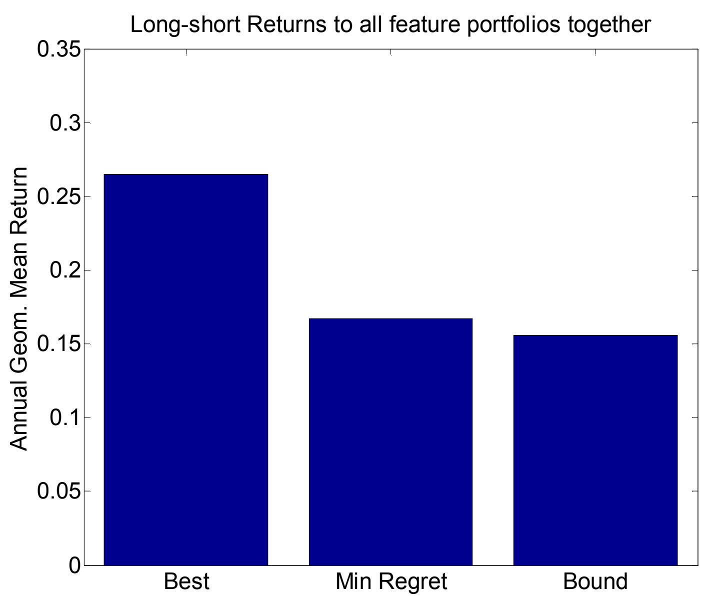

Figure 12

Annual geometric returns to extreme long-short feature portfolios conditioned on all 5 features. This creates 10 total portfolios, or 90 total long-short pairs (excluding self pairs). Best is the return of the best pair, MinRegret is the return of the generalized long-short BH strategy, and Bound is the return guaranteed to the generalized BH strategy.

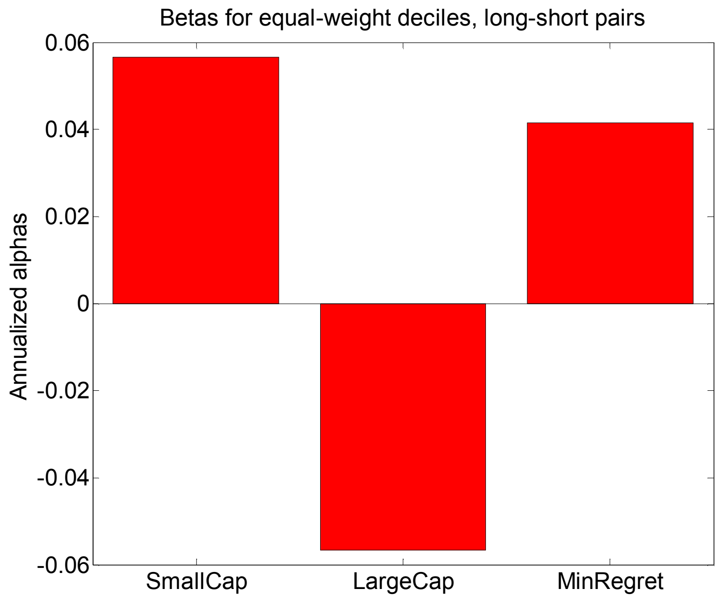

Figure 13

Risk, measured by beta, for long-short portfolios makes bets on market capitalization.

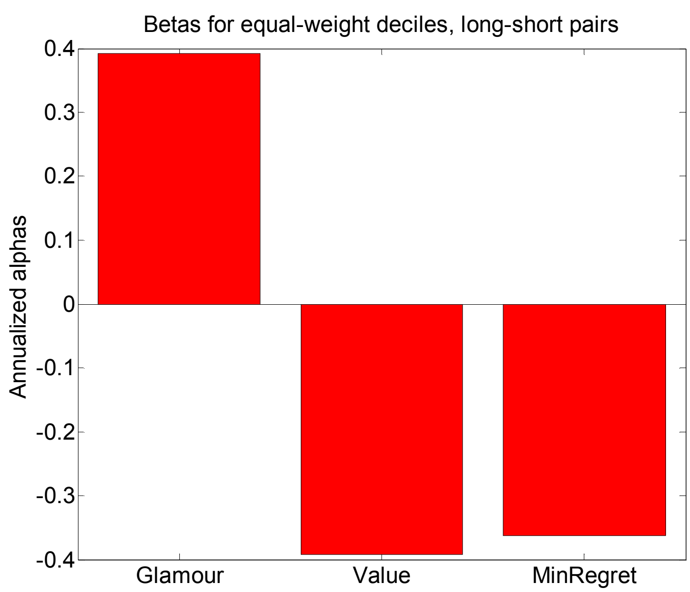

Figure 14

Risk, measured by beta, for long-short portfolios makes bets on book to market value.

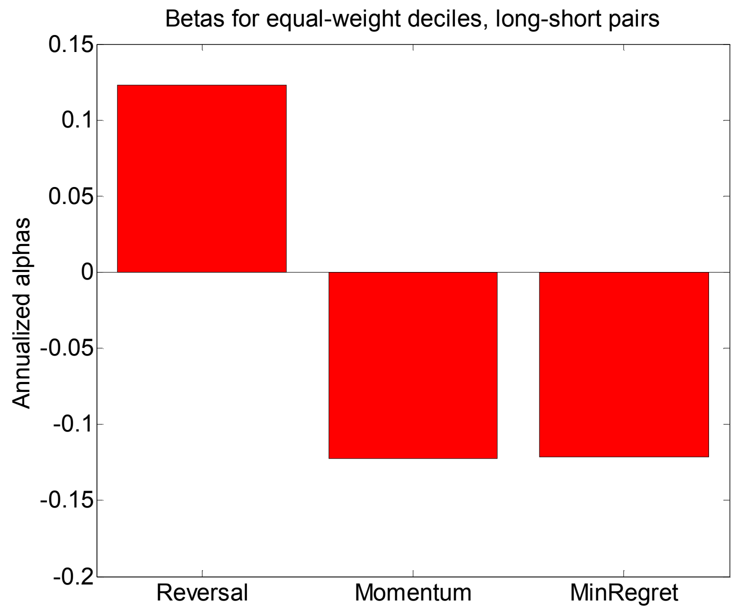

Figure 15

Risk, measured by beta, for long-short portfolios makes bets on medium-term return continuation or reversal (PAST12,2).

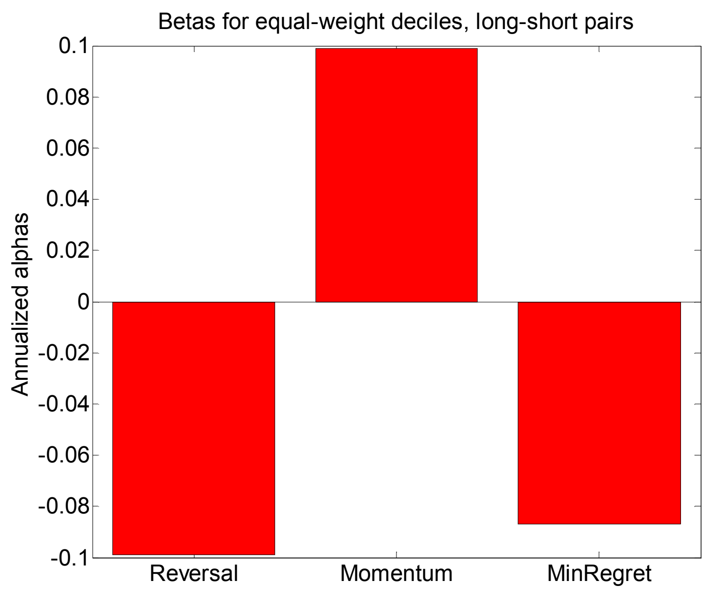

Figure 16

Risk, measured by beta, for long-short portfolios makes bets on long-term return continuation or reversal (PAST60,13).

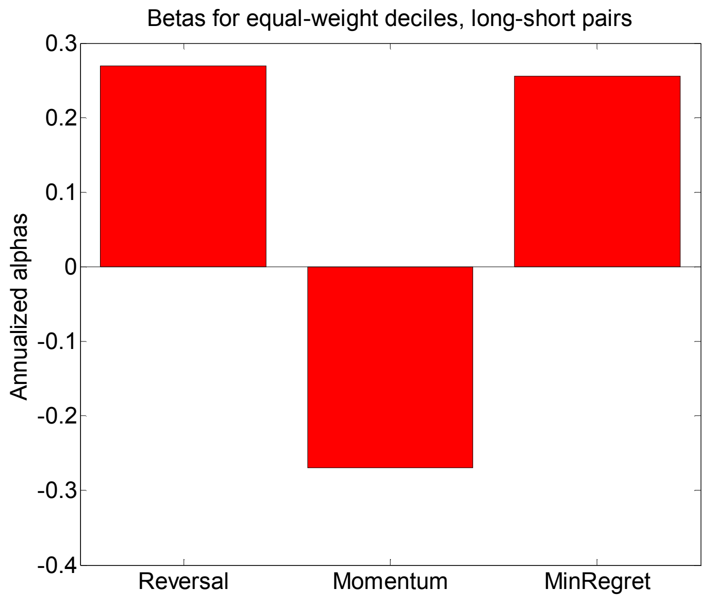

Figure 17

Risk, measured by beta, for long-short portfolios makes bets on short-term return continuation or reversal (PAST1,1).

References

Agarwal, Hazan, Kale, and Schapire (2006). “Algorithms for Portfolio Management based on the Newton Method”. Proceedings of the 23rd International Conference on Machine Learning.

Arshanapalli, Coggin, and Doukas (1998). “Multifactor Asset Pricing Analysis of International Value Investment Strategies” Journal of Portfolio Management Vol 24, No 4. p10-23

Asness (1994). Variables that Explain Stock Returns. PhD Thesis, University of Chicago 1994.

Asness (1997). The Interaction of Value and Momentum Strategies. Financial Analysts Journal March/April 1997.

Asness, Liew, and Stevens (1997). “Parallels between the cross-sectional predictability of stock and country returns.” Journal of Portfolio Management.

Asness, Friedman, Krail, and Liew (2000). “Style Timing: Value vs. Growth. Is Value Dead?” Journal of Portfolio Management. Vol 25, No 3. p50-6

Banz, Rolf W. 1981. “The Relationship between return and the market value of common stocks” Journal of Financial Economics 6, 103-126.

Basu, Sanjoy 1983. “The Relationship between earnings yield, market value, and return for NYSE common stocks: Further evidence”. Journal of Financial Economics 12, 129-156.

Black, Fischer, 1972. Capital Market Equilibrium with Restricted Borrowing. Journal of Business 45,444-455

Blum, Avrim and Kalai, Adam (1997). “Universal Portfolios with and without transaction costs”. " Proceedings of the Workshop on Computational Learning Theory

Borodin, A. El-Yaniv, R. and Gogan, V. (2003) “Can We Learn to Beat the Best Stock” Advances in Neural Information Processing Systems 16.

Borodin, A. El-Yaniv, R. and Gogan, V. (2004) “Can We Learn to Beat the Best Stock” Journal of Artificial Intelligence Research Vol 21. pp.579-594.

Branch and Chang (1985). “Tax-loss trading – is the game over or have the rules changed?” Financial Review. Vol. 20, No. 1, p55-69.

Chan, Jegadeesh, and Lakonishok (1996) “Momentum Strategies”. The Journal of Finance. Vol 51 No 5.

Cochrane, John H. (1999a). “New Facts in Finance”. Economic Perspectives Vol 23

Cochrane, John H. (1999b). “Portfolio Advice for a Multifactor World”. National Bureau of Economic Research Working Paper 710.

Conrad and Kaul (1998). “An Anatomy of Trading Strategies”. The Review of Financial Studies. Vol 11, No. 3, p.489-519.

Cover, Thomas M. (1991) “Universal Portfolios” Mathematical Finance. Vol.1 No.1 pp:1-29.

Creamer, G. and Freund, Y. (2006) “Automated Trading with Boosting and Expert Weighting” SSRN Working Paper.

De Bondt and Thaler (1985). “Does the Stock Market Overreact?” The Journal of Finance Vol. 40, No. 3. p793-805.

Fama (1998). “Market Efficiency , Long-Term Returns, and Behavioral Finance” Journal of Financial Economics Vol 49, Issue 3. p283-306.

Fama and French (1992). “The Cross-Section of Expected Stock Returns”. The Journal of Finance June 1992.

Fama and French (1993). “Common risk factors in the returns on stocks and bonds”. Journal of Financial Economics 33, 3-56

Fama and French (1996). “Multifactor Explanations of Asset Pricing Anomalies”. The Journal of Finance Vol 51, No 1,55-84.

Freund and Shapire (1995) “A decision-theoretic generalization of on-line learning and an application to boosting”. Journal of Computer and System Sciences, 55(1):119--139, August 1997.

George, Thomas J. and Hwang, Chuan-Yang (2004). “The 52-week high and momentum investing”. The Journal of Finance. Vol. 59. No 5. p2145-2176.

Gordon, Geoff (1999). “Approximate Solutions to Markov Decision Problems” PhD Thesis. CMU-CS-99-143.

Hazan, Kalai, Kale, and Agarwal (2006) “Logarithmic Regret Algorithms for Online Convex Optimization” (2006). Proceedings of 19th Annual Conference on Learning Theory.

Haussler, Kivinen, and Warmuth (1998) “Sequential prediction of individual sequences under general loss functions.” IEEE Transactions on Information Theory, 1998.

Hong and Stein (1999). “A Unified Theory of underreaction, momentum trading, and overreaction in asset markets” The Journal of Finance Vol. 54, No 6, p2143-2184.

Jegadeesh, Narasimhan (1990). “Evidence of predictable behavior of security returns”. The Journal of Finance 45 (3) 881-898

Jegadeesh, Narasimhan and Sheridan Titman. (1993) “Returns to Buying Winners and Selling Losers: Implications for Stock market Efficiency”. Journal of Finance 48,65-91

Jegadeesh, Narasimhan and Sheridan Titman (2002) “Cross-Sectional and Time-Series Determinants of Momentum Returns” The Review of Financial Studies 2002

Keim, Donald B. (2003) “The cost of trend chasing and the illusion of momentum profits”. The Rodney L White Center for Financial Research, The Wharton School 2003.

Lakonishok, Shleifer, and Vishny (1994).Contrarian Investment, extrapolation, and Risk. Journal of Finance 49, 1541-1578.

Lehmann, Bruce (1990). “Fads, Martingales, and Market Efficiency”. Quarterly Journal of Economics. Vol. 105, No.1 (February):1-28.

Lesmond, Schill, and Zhou. (2004) “The Illusory Nature of Momentum Profits”. Journal of Financial Economics 71, p349-380.

Lintner, John. 1965. “The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets” Review of Economics and Statistics 47,13-37.

Lo, Andrew (2001). “Risk Management for Hedge Funds: An Overview” Financial Analysts Journal 57,6. p16-33

Markowitz, Harry,1959. Portfolio Selection: Efficient Diversification of Investments (Wiley, New York).

Merton, R. 1987. A Simple Model of Capital Market Equilibrium with Incomplete Information”. The Journal of Finance. Vol 42, No. 3, pp:483-510.

Moskowitz and Grinblatt (1999). “Do Industries Explain Momentum?” The Journal of Finance. Vol 54, No 4, p1249-1290

Norstad, J. (2005). “Portfolio Optimization: Part 1 – Unconstrained Portfolios”. http://homepage.mac.com/j.norstad/finance/portopt1.pdf

Roll, Richard (1983). “Vas ist das?” Journal of Portfolio Management. Vol 9, No 2. P18-28

Rosenberg, Reid, and Lanstein, 1985. Persuasive Evidence of Market Inefficiency, Journal of Portfolio Management 11, 9-17

Rouwenhorst, K Geert. (1998) “International Momentum Strategies”. The Journal of Finance. Vol 53, No 1 p267-284.

Sharpe, William F. 1964 “Capital Asset Prices: A theory of market equilibrium under conditions of risk”. Journal of Finance 19,425-443