Hiding in the Hedges: Extracting Value from Well-Known Anomalies in the Equities Market☆

M.C. Creme[a], R.M. Lenain[a], J.E. Perricone[a], I.E. Shaw[a], A.J. Slottje[a]

[a] *Institute for Computational and Mathematical Engineering, Stanford University, Stanford, Calif. 94305*

Abstract

This paper examines several known abnormal return patterns in the U.S. equities markets and their use in trading strategies that attempt to anticipate price movement patterns from historical price data. We then consider the effect of news events on price movements given these patterns, as well as potential tradeability of a strategy that incorporates these signals. We investigate time-of-day anomalies across different market regimes, as well as the excess returns associated to lower-volatility stocks and momentum trading strategies.

Keywords: algorithmic trading, intraday trading, smart beta

---

1. Introduction

The investigation of anomalies in the stock market is coemergent with the practice of arbitrage and has a rich associated scholarly literature. Evidence has long persisted of market inefficiencies like the "weekend effect," [12] the "January effect," [24] and even the "Daylight Savings Time effect." [20] This is not to mention transparent violations of the law of one price like the 3Com/Palm equity carve-out and Royal Dutch-Shell price misalignments. [1] Sometimes, if transaction costs do not impede profit opportunities, arbitrageurs will eliminate these effects once they become well known. (The weekend effect no longer persists, for instance [11]). Accordingly, as the markets have become more liquid and transaction costs diminished, these effects have seen concomitant attenuation. But some biases are so ingrained, whether through institutional constellation or deep human psychological predisposition, that the effects persist. In this paper, we examine the potential implications for trading profitability when selected market anomalies are combined with signals from qualitative news data.

☆This document was produced as a final project for the class MS&E 448 and we are very grateful to Prof. Lisa Borland (Center for Financial and Risk Analytics, Stanford University) for her guidance in formulating and pursuing our project objectives.

Email addresses: `mcreme@stanford.edu` (M.C. Creme), `rlenain@stanford.edu` (R.M. Lenain), `jacobp2@stanford.edu` (J.E. Perricone), `ieshaw@stanford.edu` (I.E. Shaw), `slottje@stanford.edu` (A.J. Slottje)

One of the anomalies we consider involves momentum effects, instantiated on a variety of timescales. Asset prices have positive observed serial correlations, which have been shown to be induced by investor behavioral bias and market risk structure. It has been shown that different investors react to the arrival of new information at a different speed. Typically, after initial under-reaction to news, investors extrapolate past behavior and create positive price momentum. As market information settles, the price movements caused by the momentum effect exhibit reversion to the efficient price. [16, 15] As a result, all momentum strategies display significant fat tails, a property relating to crash risk, with equity momentum exhibiting the highest tail risk and drawdowns. We develop a momentum strategy that aims to mitigate this excess kurtosis by overlaying momentum strategies on different scales.

Previous work on longer-horizon momentum strategies includes that of Jegadeesh and Titman, [18] who determined the profitability of buying stocks with high returns and selling stocks with low returns on a semiannual timescale. Persistence and robustness to data-mining of these effects has been demonstrated. [19] Additionally, Gao et al. have shown a correlation in intraday returns that would statistically corrobate the success of such a strategy on a daily timescale. [14] In particular, they deduce existence of price movement patterns consistent with shifts in observed asset prices due to market-maker bookkeeping; returns from the first 30 minutes of the trading day are found to predict the last 30 minutes. In other studies, overnight returns are found to inversely predict next-day returns. [5] Of potential interest for our sentiment analysis, Gao et al. find an association between macroeconomic news and strength of intraday momentum effects.

On the other hand, owing to transaction costs, profitability does not immediately follow from these findings of statistical correlation. It has been generally found that transaction costs obviate the benefit to trading on systematic intraday statistical correlation in prices. [21] For example, using slightly more complicated holding rules, Lam et al. find transaction costs to obviate profitability of certain intraday momentum strategies. [22] Because of the holding requirements of a first 30/last 30 strategy (*i.e.*, holding only for the last 30 minutes), liquidity issues may pose a structural obstacle to traders hoping to turn a profit; Branch and Ma hypothesize that such profits may primarily accrue to market-makers. [5] This would be consistent with other observed failures of arbitrage. [25] We find the correlations not to exist at all, and that the strategy is unprofitable by itself, with the exception of certain market regimes; we accordingly focus on high volatility and incorporating outside signals like volume and volatility as a means to profitability (which we attain, although not at excess to the market). This is consistent with the results of Gao et al. that momentum effects are most pronounced when volatility is high. [14]

We additionally consider the well-known anomaly that stocks with low volatility outperform those with high volatility. Ciliberti et al. demonstrate existence of this effect and attribute it to investor underassessment of dividend yields, since lower-volatility stocks tend to pay higher dividends but do not support the heuristic hold that "glittering" stocks can maintain on investors. [8]

Additional, less obvious, factors contribute to the profitability of this strategy as well: institutional investor bias, market bias, and a 'compounding effect.' [8] The bias of institutional investors may be attributed to compensation structure; these investors tend to focus on high volatility stocks whose short-term upswings might be useful to pad their bonuses. This leads to added volatility and overvaluation. The resulting propensity of retail and institutional investors toward hyper-focus on consequently even higher-volatility stocks entrenches these effects.

Now consider an asset that is highly volatile. One day it increases 20% in value, while the next day it decreases 20% in value. After these two volatile days, though the returns sum to 0, the asset is actually 4% down in value. This is the compounding effect. In concordance with these dynamics, we find substantial returns to a low-volatility strategy, in support of the analysis of Ciliberti et al.

It has additionally been found that strategically combining momentum trading and low-volatility trading can boost returns. [23] Many of the signals deployed in such strategies map information to state spaces increasingly liminal to price equilibrium; locating a strategy at the intersection of these signals represents an opportunity that is as yet possibly more robust to arbitrage. We develop a combined-signals strategy which obtains 235.9% returns from January 2005 to June 2017, a return of about 150% of the S&P 500's.

Building on these results, we attempt to cultivate our investor sentiment signals data in order to incorporate it into the trading strategies we have refined. We use a dataset that indicates direction and existence of sentiment and examine whether it is possible to trade on these signals before the market fully adjusts. The literature on sentiment signaling has met with mixed results in proving a correlation between news data and stock returns, but Boudoukh et al. show that results may be obtained by constraining the news dataset to control relevance of news. They find that news days are associated with continuation of stock price patterns, while normal days show reversals. On the other hand, news days are also associated with greater volatility and extreme returns. [4] Moreover, knowledge of the extended cumulation of excess returns following a news event has been known since at least Bernard and Thomas's study of post-earnings-announcement drift. [2] Several studies have found these dynamics most pronounced among smaller stocks, and especially following negative news. [7, 16, 15] One may interpolate time dependence in these patterns. For example, managers may look to capitalize on limited investor attention by releasing unfavorable news on Fridays, resulting in negative abnormal returns on Fridays. [9] Indeed, negative news on Fridays sees delayed response at attenuated volume. [10]

We develop a machine-learning algorithm to take advantage of this dataset and these systematically ex-post present price adjustments. However, predictability is not very strong, because we do not have access to the requisite data to train our algorithm to detect either immediate or long-term price moves. This represents an important area for future research.

We hypothesize that we might improve our intraday momentum strategy by putting to use the high volatility and tendency toward continuation exhibited by the market on news days. It has also been shown in the aggregate that news data provide a way to anticipate volatility ahead of the VIX, potentially enhancing our low-volatility and volatility-dependent momentum strategy results. [3] We find that our final, combined-signals strategy performs worst during periods of low volatility and weak-to-none market direction. To motivate our work in both sentiment signaling and trading algorithm development, we briefly discuss the idea of using news data to signal these periods, thereby enabling dynamic weighting of our combined-signals strategy in order to adapt to unfavorable regimes.

2. Data

2.1. *Preliminary exploration of the intraday returns space: Bloomberg and Google Finance*

We began our investigation of the previously discussed return anomalies with a limited dataset in order to acquire intuition for marketplace price evolution dynamics. Specifically, we looked at publicly available historical price data from Google Finance as well as a six-month returns series available through Bloomberg. The prices we pulled were for the SPY and other similar exchange-traded funds. We performed least-squares regressions to test association of the return in the last half hour of the data with the first half hour in the day. We additionally performed similar tests to ascertain correlation of the daily return with the intraday return.

It is entirely possible the small size of the dataset limited the significance of our results, but initial tests were not encouraging. Both correlation coefficients and significance were identically zero with high precision, and tests with the more extensive Quantopian data yielded similar results. We discuss these further in Section 3.

2.2. *Quantopian data*

Quantopian provides an environment for algorithmic backtesting. Using the proprietary Pipeline API available through Quantopian, users can develop strategies that regularly rebalance a portfolio of stocks using rules based on trailing-window criteria, with timing determined by the user (for example, in our first 30/last 30 strategy, we set the rebalancing time so as to purchase stocks daily at 3:30 and additionally require liquidation at 4 o'clock.) In our implementations, we limited the universe to different tradeable sets of stocks defined in the Quantopian API because of the high liquidity demands of our strategies.

2.3. *Inferess news data*

Our qualitative data came from Inferess. The Inferess dataset features the relevant company's ticker, the date and time of the article release as well as different measures of sentiment about the article. These measures include an estimation of the probability of the direction of the news item (positive or negative) along with the counts of positive and negative words throughout the article. In order to find matching stock prices for the date and time associated to each article, daily stock prices were pulled from Google Finance. We tried to get minute-level data from several sources (Google Finance and Bloomberg), but limitations on the publicly available datasets to within a recent period in each case prevented us from using them in conjunction with the Inferess data, which were observed at an earlier time.

3. Methods

3.1. Momentum strategy

3.1.1. Monthly-horizon momentum

We describe a long-only long-term momentum strategy, which we find to successfully mitigate the effects of the fat tails described above (albeit with limited upside). The strategy comprises a long-only momentum trade that rebalances monthly. Specifically, we define our initial asset universe $\boldsymbol{S}$ as the top 500 most actively traded stocks over the past year. At the opening of each month, the historic 200-day trailing return is calculated. The universe is then sorted in decreasing order. A hyper-parameter of the number of assets to include in the monthly momentum strategy, N_{ml}, filters the asset universe, so that only the top N_{ml} stocks are selected. For each stock within this restricted subset, if the current daily return is greater than the 200-day moving average, a rank-weighted long position is entered in the stock. This process is repeated monthly, resetting all current positions.

Algorithm 1: Long-Only Monthly Momentum Strategy

```
Reset all held positions; i.e., set w_{s_i} = 0;
S_sorted ← initialize and sort asset universe according to 200-day trailing return ;
/* Define the momentum list of top momentum stocks */
M = S_sorted[0 : N_ml];
f_h = Σ_{j=1}^{N_ml} j ;
for s_i ∈ M do
    if r_{s_i}^{(t)} > 1.005 * MAVG(r_{s_i}) then
        /* Purchase weight s_i in security */
        w_{s_i} = rank(M(s_i)) / f_h;
```

The hyperparameter N_{ml} is chosen to be 30, but can be optimized in backtesting.

3.1.2. Intraday momentum

As discussed in section 1, intraday price movements have been shown to exhibit a degree of predictability. One explanation of this phenomenon relates to traders' execution behavior. Institutional traders are often judged based on how they purchase an asset as compared to its volume weighted average price (VWAP) throughout the day. Consequently, traders often aim to execute during periods of high volume. If news at the opening impacts an asset price, rational traders will purchase a fraction of their required goal at the beginning of the day, when volume is high and the VWAP is not yet set. Throughout the day, the trader purchases smaller amounts waiting for the VWAP to settle. Finally, at the close (when trading volume spikes), the trader finishes executing. If the news is good, this would cause the trader to put buying pressure on the asset during the beginning of the day and then on the end of the day. Similar results obtain for purchases as for sales. [6]

Let us mathematically detail such a strategy. On trading day t, break the day into $t_1, \ldots, t_i, \ldots t_{13}$ segments corresponding to each thirty-minute window throughout the training day. We write the return in window i of day t as $r_i^{(t)} = \frac{p_{t_{i+1}} - p_{t_i}}{p_{t_i}}$, with p_{t_i} denoting the price of the asset at the end of period i.[1] Recall the results discussed above; under appropriate market conditions, the returns during the first period and the penultimate period have a high correlation on the closing periods returns. Let $A = \{a_1, \ldots, a_i\}$ denote the asset universe in the scope of our strategy, and let L_{im} denote the lookback used by the strategy. Since the intraday strategy aims to take advantage of tail events on the intraday scale, the algorithm uses historic intraday volatility and volume to define execution thresholds.

In particular, for each period i of the trading day, the algorithm calculates the volatility and volume that occurred in period i for a lookback period $t-1, t-2, \ldots, t-L_{im}$. For an individual asset, let $\boldsymbol{V}_i = \{v_i^{(t-1)}, v_i^{(t-2)}, \ldots, v_i^{(t-L_{im})}\}$ denote the volume traded during window i of each day in the lookback period.[2] Calculate the same lookback for $\boldsymbol{\sigma}_i, \boldsymbol{R}_i$, denoting the historic intraday volatility and return. On each day t at the onset of the last period of the day, t_{13}[3]:

[1] p_{t_0} is defined as the closing price of the previous day.

[2] $v_i^{(t-1)}$ denotes the volume traded during window i on day $t-1$.

[3] 30 minute period ending at market close.

Algorithm 2: Intraday momentum strategy

```
Initialize γ_i^vol, γ_i^volume as quantile thresholds for period i for volatility and
  volume respectively ;
/* At the beginning of period i = 13                                          */
for a_i ∈ A do
    /* Define trigger periods                                                 */
    Let I = {i = 1, 12} ;
    /* Fetch Historic Intraday data                                           */
    Calculate V_i, σ_i, R_i   ∀  i ∈ I ;
    Calculate v_i^(t), σ_i^(t), r_i^(t)   ∀  i ∈ I;
    /* Let F̂^{-1}(x, X) be the empirical quantile of an observed quantity x
       with respect to the data X.                                            */
    if F̂^{-1}(v_i^(t), V_i) > γ_i^volume and F̂^{-1}(σ_i^(t), σ_i) > γ_i^vol   ∀  i ∈ I then
        /* If return in first and penultimate period are greater than
           thresholds                                                         */
        if r_i^(t) > 0 and F^{-1}(r_i^(t), R_i) > γ_i^{r_buy}   ∀  i ∈ I then
            /* order amount w_{a_i}^buy                                       */
            order(w_{a_i}^buy);
        if r_i^(t) < 0 and F^{-1}(r_i^(t), R_i) < γ_i^{r_sell}   ∀  i ∈ I then
            /* order amount w_{a_i}^sell                                      */
            order(w_{a_i}^sell);
/* At end of day close out position                                           */
for a_i ∈ A do
    if |w_{a_i}| > 0 then
        order(−w_{a_i})
```

The intraday momentum strategy contains the hyperparameters

$$\boldsymbol{\theta} = \{w_{a_i}^{sell}, w_{a_i}^{buy}, \gamma_i^{r_{buy}}, \gamma_i^{r_{sell}}, \gamma_i^{volume}, \gamma_i^{vol}\}$$

Since the intraday strategy makes use of what has been shown to be a superior correlation in times of high volatility and high volume [13], the parameters $\gamma_i^{volume}\gamma_i^{vol}$ are chosen to be .9. If the volume and volatility criterion are met, the algorithm checks for a buy (sell) signal. If the returns are positive (negative) and greater than (less than) $\gamma_i^{r_{buy}} = .75$ ($\gamma_i^{r_{sell}} = .1$),[4] a buy (sell) order is placed on the stock. When the algorithm is run in conjunction with the long term momentum strategy,[5] the proportion w_{a_i} is set to be a fixed percentage of the absolute value of the holdings of the asset, where $w_{a_i}^{buy}$ is a positive percentage and $w_{a_i}^{sell}$ is negative.

[4] The historic returns are filtered for those greater than (less than) zero before the quantile function is applied.

[5] This is how it is run in our experimental portfolio.

3.1.3. Discussion

We develop a strategy which overlays the long-only monthly momentum strategy and intraday momentum strategy. We restrict the asset universe $\boldsymbol{A}$ of the intraday strategy to include only assets held by the long-only monthly momentum portfolio. The intuition is that the monthly momentum strategy captures a long-term risk factor (which is susceptible to large downswings), while the intraday strategy captures profit by betting on intraday anomalies, trader behavioral activity, and the market microstructure.

3.2. Low-volatility strategy

3.2.1. Algorithm

We turn now to the low-volatility strategy. Low-volatility strategies, eponymously embodied, seek to short high-volatility stocks and to go long on low-volatility stocks, taking advantage of structural factors that favor long-term returns from equities associated with lower volatility. [8] Our implementation of the strategy takes the trailing 100-day historical volatility of the universe, ranks the equities by volatility, and shorts the highest volatility assets while taking a long position on the lowest volatility assets.

For this strategy we set our universe U_{lv} to be the Q1500 universe defined by the Quantopian API, which is the 1500 most tradeable assets by 200-day average dollar volume, capped at 30% of equities allocated to any single sector. This is necessary to maximize the possibility that orders may be filled within the day set for rebalancing. We decided to go short and long on 100 equities each ($n_h = n_l = 100$) and weight them proportional to historical risk.

Algorithm 3: Low Volatility Rebalancing Algorithm

$P_c =$ portfolio value outside of strategy ;
$n_u = |U_{lv}|$;
$\Sigma_{100} = \{\sigma_{100}(s_i) : \forall s_i \in U_{lv}\}$;
$f_h = \sum_{j=1}^{n_h} j$;
$f_l = \sum_{k=1}^{n_l} k$;
for $s_i \in U_{lv}$ **do**
  if $rank(\sigma_{100}(s_i)) \geq n_h$ **then**
    $p_i = -1 \cdot \frac{rank(\sigma_{100}(s_i))}{f_h} \cdot P_c$;
  if $rank(\sigma_{100}(s_i)) \leq n_u - n_l$ **then**
    $p_i = \frac{rank(\sigma_{100}(s_i))}{f_l} \cdot P_c$;
  else
    $p_i = 0$

3.2.2. Discussion

The long side of this strategy seeks to take advantage of the steady (but historically undervalued) profitability of low-volatility equities discussed above, associated with blue-chip companies like Johnson & Johnson or Lockheed Martin. In order to establish beta-neutrality, the strategy has a short side as well.

The added benefit of pairing short and long sides is that the strategy becomes self-financing. The strategy is therefore designed as a small, consistent profit strategy due to its small margins, beta neutrality, and medium frequency.

3.2.3. Selection of the rebalancing horizon

We explored daily, weekly, and monthly horizons, choosing a suitable long/short rebalancing to maintain beta neutrality. (We determined this empirically; another suitable, and empirically more or less equivalent, strategy would be dollar neutrality.) After testing, monthly rebalancing generated superior returns, along with improved measures of maximum drawdown, Sharpe ratio, and beta. We believe that the longer horizon in monthly rebalancing capitalizes on the compounding effect of the short side. If the horizon is too short, there is a risk of being caught in the price upswing of a volatile asset and being forced to close that short position unprofitably. A longer time horizon allows the swings to offset the negative position, as they compound over time.

3.3. Combined strategy

So far our strategies have explored two long-term risk factors, momentum and volatility. The former captures profit from serial autocorrelation of returns whereas the latter exploits the reversion trend in high-volatility/low-volatility assets. The two strategies therefore capture different risk premia within the market and should exhibit low correlation. The low-volatility strategy performs well during market downturns, whereas the momentum strategy takes greater advantage of upswings.

Amidst these long-horizon dynamics, intraday momentum trading represents a source of uncorrelated alpha that simply bets on intraday anomalies. In an effort to strengthen our portfolio's resilience to downturns and bolster our signals, we accordingly combine these three different strategies. The monthly-horizon momentum strategy and low-volatility strategy act on different subsets of the asset universe, whereas the intraday strategy acts on all assets that have holdings in the portfolio. For simplicity, let us detail some notation. First, let $S_{lv}(A, w)$ be the low-volatility strategy called on asset universe A given weighting w. Let A_{lv} denote the positions the low-volatility strategy takes. Similarly define $S_{ml}(A, w)$ as the momentum strategy called on with the asset universe A with portfolio weight w. Let A_{wl} be the positions the monthly horizon momentum strategy takes. We use the same notational definitions here as in the intraday strategy above.

Algorithm 4: Combined strategy

```
A ← Top 500 traded by volume stocks;
/* On Each Month                                   */
S_ml(A, w_ml);
A_ml ← S_ml(A, w_ml).A_ml;
if mid-month then
    S_lv(A A_ml, w_lv);
    A_lv ← S_lv.A_lv;
for each day within the month do
    A_intraday(A_ml ∪ A_lv);
```

Now, since the low-volatility strategy is self-financing, the weights w_{ml}, w_{lv} are critical in the strategy's performance. If $w_{ml} > 1$ (*i.e.*, levered momentum), one would expect more volatility in the return series. Conversely, if the low-volatility strategy dominates, the return series should exhibit better resilience in downturns. We discuss these results below.

4. Momentum strategy results

The hyper-parameters $\gamma_i^{r_{buy}}, \gamma_i^{r_{sell}}, \gamma_i^{volume}, \gamma_i^{vol}$ were chosen based on intuition, not optimization, and can therefore be further improved. Since we are looking for extreme tail events, $\gamma_i^{volume} = \gamma_i^{vol} = .9$. The quantile buy threshold $\gamma_i^{r_{buy}} = .75$ whereas the sell threshold $\gamma_i^{r_{sell}} = .1$. The results of taking an intraday position in $\frac{1}{2}, \frac{1}{5}, \frac{1}{10}$ of the current allocation are shown below. The top performing strategy, a 10% allocation, yielded a return of 105.6%, falling below the benchmark, and exhibited a large maximum drawdown of -59%. We find 7% alpha, with a fairly low Sharpe ratio of .35. During the financial crisis the strategy's previously high excess returns mitigated losses. Furthermore, when compared to the long-horizon momentum strategy, the overlaid results show a stronger resilience to downturn. More work is needed for the strategy to perform at par with the market, but the signals we have extracted appear to be meaningful.

Figure 1: Overlaid Momentum Strategy Results

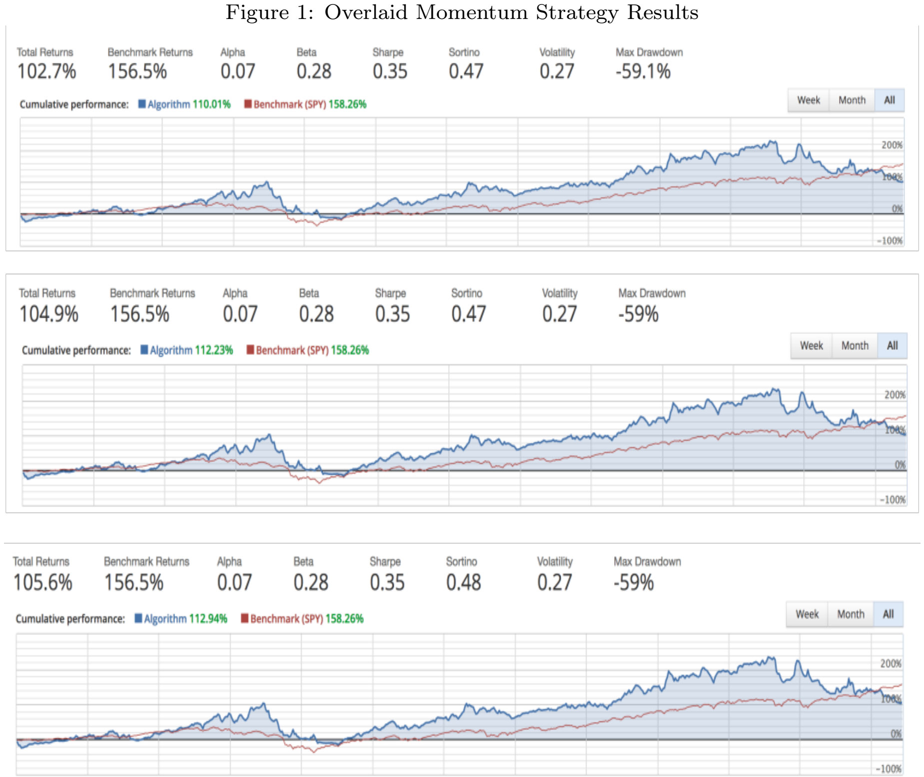

% Allocated	Strategy Returns	Benchmark Returns	Alpha	Beta	Sharpe	Sortino	Volatility	Max Drawdown
50 position	102.7 %	156.5 %	.007	.28	.35	.47	.27	-59.1 %
20 position	104.9 %	156.5 %	.007	.28	.35	.47	.27	-59 %
10 position	105.6 %	156.5 %	.007	.28	.35	.48	.27	-59 %

Further improvements can be achieved in optimizing the hyper-parameters of the strategy.

5. Low-volatility strategy results

The low-volatility strategy with monthly rebalancing exhibited performance of 23.8% returns against 146.1% benchmark returns on the SPY during the period from June 2005 to June 2017. The Sharpe ratio of 0.26 is comparable to the results of Ciliberti et al., who find a Sharpe ratio of 0.24. [8] To its benefit, the strategy has a low β of 0.05, as desired.

Figure 2: Low-Vol Strategy Results

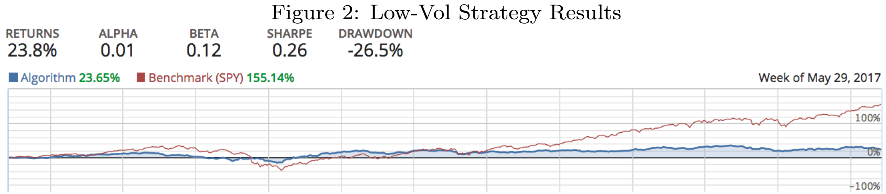

6. Combined strategy results

The combined strategy's returns proved significantly higher than the results from the individual strategies pursued alone, suggesting a benefit to signal combination. Holding 95% of the portfolio in the monthly horizon momentum strategy, 95% in the self-financing low-volatility strategy, and betting 50% of the position using our intraday strategy,[6] we achieved an excess return over the market of $\approx 125\%$, with a Sharpe ratio of .53 and 12% alpha. Nevertheless the drawdown, though diminished in comparison to our momentum-only strategy, is still high at -51.9%. In an effort to reduce our drawdown, we backtested the combined strategy with a limit of 50% on the portfolio apportionment to the momentum strategy. This improved results, with the resulting Sharpe ratio increasing to .64, beta decreasing to .17, volatility nearly halving to 15%, and the maximum drawdown increasing to -27.2%. The results are shown below.

Figure 3: Combined Strategy Results

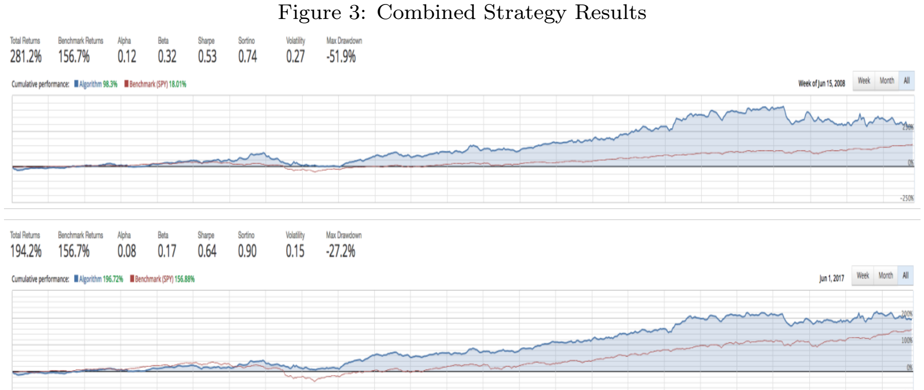

Momentum Weighting	Strategy Returns	Benchmark Returns	Alpha	Beta	Sharpe	Sortino	Volatility	Max Drawdown
95% position	281.2 %	156.7 %	.12	.32	.53	.74	.27	-51.9 %
50% position	194.2 %	156.7 %	.008	.17	.64	.90	.15	-27.2 %

[6] We set the quantile return thresholds to .9 and .1, respectively.

7. Sentiment signaling results

7.0.1. Features and labels

We employed several machine learning algorithms in an attempt to classify the observations in our news sentiment dataset. In performing this classification, parts of the Inferess dataset were used as features: the positive and negative sentiment, the percentage of total positive and negative words, and the count of positive and negative sentences. We first attempted binary classification with labels indicating whether or not the underlying stock rose by more than 0.5% in the day following the article's release. We also implemented multi-class classification with labels indicating a 0.5% percent rise in stock price, 0.5% percent fall in stock price, and neither of these events. In order to visualize the data, the data were properly normalized and principal components analysis was performed for each classification.

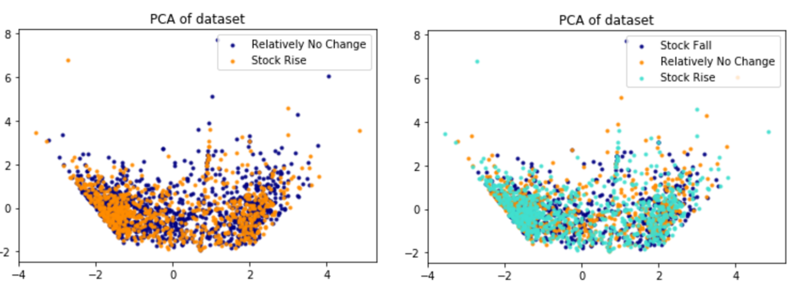

7.0.2. Results

Using the features and labels described, we attempted many different machine learning algorithms, including support vector machines with multiple different kernels such as polynomial, linear, and radial basis function kernels. In addition, neural networks, stochastic gradient descent, linear regression, and gradient boosting were all pursued in an effort to classify the news data. The results were of limited significance: most of the classifiers classified the majority of the test set under one label. Here are the results when using the neural network in the binary classification problem:

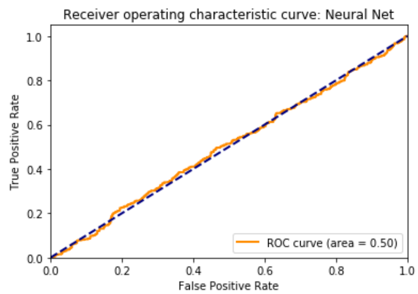

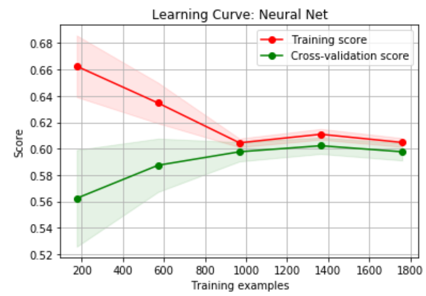

	precision	recall	f1-score	support
0.0	0.60	0.95	0.74	578
1.0	0.33	0.04	0.07	378
avg / total	0.50	0.59	0.47	956

Most of the other classifiers yielded very similar results. As there was no statistical correlation found between the features from the Inferess dataset and daily stock price changes, we were unable to augment the other signals found with this news signal.

8. Biases in results

We considered our strategies successful if they exceeded a rather low risk benchmark (as opposed to if they simply obtained positive returns). We also used the SPY index as a comparison, which created bias in our approach toward beating the market instead of generating consistent returns. (It is, of course, not always easy to generate consistent returns when the market is doing poorly.)

An important point to note is that we backtested strategies on the horizon 2003 - 2017, in order to explore our strategies' performance during the subprime crisis, as well as the pre-crisis and post-crisis eras. This allowed us to try to optimize our hedging for all types of market regimes. However, we are possibly overfit to the market of the early 21^{st} century in ways we do not foresee. Testing with randomly generated data would be a good step toward ensuring this strategy is as robust to tail-end events as the theory would predict, and not just the financial crisis.

One problem we encountered when running our first trading algorithms was the lack of liquidity in certain stocks, particularly when running intraday strategies which required closing out all of our positions in the last thirty minutes of each day. To sidestep such issues, we decided to limit ourselves to highly traded stocks. This creates bias in our data, as we are likely overfit to a highly liquid universe in U.S. equities.

It also goes without saying that we are likely overfit to the U.S. market in particular of (highly traded) stocks. This limitation was necessary in order for us to access data useful for our simulations, but we were not able to cross-validate on, say, illiquid stocks or the Japanese market, so we do not know if our results obtain in such instances. (By way of illustration, in the 1980s analysis of the "weekend effect," it was found that Japan and Australia also experienced Monday effects - on Tuesday. [17])

9. Analysis

We find the best returns to accrue to the portfolio that makes use of all the types of signals and strategies discussed - in fact, it is the only portfolio to outperform the market. This suggests that it is possible to combine risk factors for excess return even when the factors do not generate profits in isolation. In particular, the example of predictable intraday pricing dynamics is instructive. While trading solely on such signals intuitively (and in our results, empirically) should run up against challenges to arbitrage (otherwise prices would evolve differently), we were able to form a profitable strategy by combining information evidently not considered in the larger market, possibly due to the differing horizons on such information. It may well be institutional, with trading desks not reconciling the information in one place, or it may be psychological, with sophisticated investors classifying different anomalies into multiple categories without considering their combination. Although the market has arbitraged away profits to strategies like intraday and multi-day momentum on their own, perhaps the intersection is more than the sum of its parts.

10. Future work

We see two broadly defined areas for future work building on our results. First, there is a need for cross-validation of our results in other asset classes and international markets, in order to guard against data mining. We did not develop these results ourselves owing to constraints on the availability of the requisite data.

Secondly, within the strategy we have already developed, we have not fully optimized hyperparameters (when we ought to trade during the day, optimal limits on leverage). In particular, we find significant improvement on a risk/reward basis from changing the weighting between strategies in our multi-strategy portfolio. Accordingly, we wish to develop a method to track profits and losses to each strategy, in order to empirically guide and validate optimization of portfolio weights.

Moreover, these changes suggest larger ones: *i.e.*, dynamically weighting the strategies. For example, if the momentum strategy exhibits diminishing returns, one might reallocate a percentage of the portfolio to the low-volatility strategy. This will help mitigate the risk associated with the strategies during high-volatility regimes, and takes advantage of the complementary nature of the underlying risk factors. Expanding to other known market anomalies in the "smart beta" universe provides an avenue to generalize this approach as well.

We see an exciting opportunity for news data to guide these strategic questions as a means to signal extraction. We believe that combining sentiment signals with minute-level data or longer daily price series may produce a more significantly predictive signal than we were able to obtain. While limitations on our access to such data precluded us from fully using sentiment signals to our advantage, the literature on news sentiment suggests that an effective machine-learning algorithm trained on sentiment data could well complement our intraday strategy as a short-horizon alpha boost, in conjunction with a dynamic rebalancing strategy.

11. References

[1] Nicholas Barberis and Richard Thaler. *Handbook of the Economics of Finance*, volume 1, chapter 18, pages 1053–1128. Elsevier, 2003.

[2] Victor L. Bernard and Jacob K. Thomas. Post-earnings-announcement drift: Delayed price response or risk premium? *Journal of Accounting Research*, 27(Current Studies on The Information Content of Accounting Earnings):1–36, 1989.

[3] Svetlana Borovkova, Evgeny Garmaev, Philipp Lammers, and Jordi Rustige. Sensr: A sentiment-based systemic risk indicator. *De Nederlandsche Bank Working Paper*, (553), 2017.

[4] Jacob Boudoukh, Ronen Feldman, Shimon Kogan, and Matthew Richardson. Which news moves stock prices? a textual analysis. *NBER Working Paper*, (18725), January 2012.

[5] Ben Branch and Aixin Ma. Overnight return, the invisible hand between intraday returns? *Journal of Applied Finance*, 22(2), 2012.

[6] Christian T. Brownlees, Fabrizio Cipollini, and Giampiero M. Gallo. Intradaily volume modeling and prediction for algorithmic trading. *Journal of Financial Econometrics*, 9(3):489–518, 2011.

[7] Wesley S. Chan. Stock price reaction to news and no-news: drift and reversal after headlines. *Journal of Financial Economics*, 70(2):223–260, November 2003.

[8] S. Ciliberti, Y. Lempérière, A. Beveratos, G. Simon, L. Laloux, M. Potters, and J.-P. Bouchaud. Deconstructing the low-vol anomaly.

[9] Aswath Damodaran. The weekend effect in information releases: A study of earnings and dividend announcements. *Review of Financial Studies*, 2 (4):607–623, 1989.

[10] Stefano Dellavigna and Joshua Pollet. Investor inattention and friday earnings announcements. *Journal of Finance*, 64(2):709–749, 2009.

[11] Peter Fortune. Weekends can be rough: Revisiting the weekend effect in stock prices. *Federal Reserve Bank of Boston Working Paper*, 98(6), September 1998.

[12] Kenneth French. Stock returns and the weekend effect. *Journal of Financial Economics*, 8(1):55–69, 1980.

[13] Lei Gao, Yufeng Han, Sophia Zhengzi Li, and Guofu Zhou. Intraday momentum: The first half-hour return predicts the last half-hour return, 2014.

[14] Lei Gao, Yufeng Han, Sophia Zhengzi Li, and Guofu Zhou. Market intraday momentum, 2017.

[15] Harrison Hong and Jeremy C. Stein. A unified theory of underreaction, momentum trading and overreaction in asset markets. *Journal of Finance*, 56(6):2143–2184, December 1999.

[16] Harrison Hong, Terence Lim, and Jeremy C. Stein. Bad news travels slowly: size, analyst coverage, and the profitability of momentum strategies. *Journal of Finance*, 55(1):265–295, February 2000.

[17] Jeffrey Jaffe and Randolph Westerfield. The week-end effect in common stock returns: The international evidence. *Journal of Financial Econometrics*, 40(2):433–454, 1985.

[18] Narasimhan Jegadeesh and Sheridan Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, 48(1):65–91, 1993.

[19] Narasimhan Jegadeesh and Sheridan Titman. Profitability of momentum strategies: An evaluation of alternative explanations. *Journal of Finance*, 56(2):699–720, 2001.

[20] Mark Kamstra, Lisa Kramer, and Maurice Levi. Losing sleep at the market: The daylight saving anomaly. *American Economic Review*, 90(4):1005–1011, 2000.

[21] Oleg Komarov. Intra-day momentum, January 2017.

[22] Vincent Wing-Shing Lam, Terence Tai-Leung Chonga, and Wing-Keung Wong. Profitability of intraday and interday momentum strategies. *Applied Economics Letters*, 14(15):1103–1108, 2007.

[23] Ari Polychronopoulos. Building a better beta: Combining fundamentals weighting, low volatility, and momentum strategies. *Research Affiliates White Paper*, October 2014.

[24] Michael S. Rozeff and William Kinney. Capital market seasonality: The case of stock returns. *Journal of Financial Economics*, 3(4):379–402, 1976.

[25] Andrei Shleifer and Robert W. Vishny. The limits of arbitrage. *Journal of Finance*, 52(1):35–55, March 1997.