



[Finance Program]

“Multi-Armed Bandit Approach to Portfolio Choice Problem”

Guneykan Ozkaya, Yaping Wang

Director: Filippo Ippolito

10/07/2020

ABSTRACT IN ENGLISH (100 words):

In this paper, we combine several Multi-Armed bandit algorithms with methodologies from finance literature and apply it to portfolio choice problem. Our results show that when we combine bandit algorithms with methodologies that take account of the non-normal distribution of returns, portfolio performance improves. Our results show that if contextual bandit algorithms applied to portfolio choice problem, given enough context information about the financial environment, they can consistently obtain higher Sharpe ratios compared to classical methodologies, which translates to fully automated portfolio allocation framework.

ABSTRACT IN CATALAN/ SPANISH (100 words)

En este documento, combinamos varios Multi-Armed bandit algorithms con metodologías de la literatura financiera y lo aplicamos al problema de elección de cartera. Nuestros resultados muestran que cuando combinamos estos algoritmos con metodologías que tienen en cuenta la distribución no normal de los rendimientos, el rendimiento de la cartera mejora. Adicionalmente, dada la suficiente información sobre el contexto del entorno financiero, nuestra investigación muestra que si los denominados contextual bandit algorithms se aplican al problema de elección de cartera, estos pueden obtener consistentemente mayores Sharpe Ratio en comparación con las metodologías clásicas, lo que se traduce en un marco de asignación de cartera totalmente automatizado.

KEYWORDS IN ENGLISH (3):

UCB1, Thompson Sampling, Probabilistic Sharpe Ratio

KEYWORDS IN CATALAN/ SPANISH (3):

UCB1, Thompson Sampling, Probabilistic Sharpe Ratio



Multi-Armed Bandit Approach to Portfolio Choice Problem

of Economics

Yaping Wang

Msc. in Finance

Barcelona Graduate School
of Economics

G  neykan Ozkaya

Msc. in Finance

Barcelona Graduate School

Advisor: Jos   Su  arez Lled  o

Abstract

In this paper, we combine several Multi-Armed bandit algorithms with methodologies from finance literature and apply it to portfolio choice problem. Our results show that when we combine bandit algorithms with methodologies that take account of the non-normal distribution of returns, portfolio performance improves. Our results show that if contextual bandit algorithms applied to portfolio choice problem, given enough context information about the financial environment, they can consistently obtain higher Sharpe ratios

compared to classical methodologies, which translates to fully automated portfolio allocation framework.¹

Keywords – UCB1, Thompson Sampling, Probabilistic Sharpe Ratio

¹Implementation of all algorithms can be found at <https://github.com/guney1>

Contents

Contents

1 Introduction 1 2 Literature Review 2

3 Methodology	4 3.1 Orthogonal Portfolios	4 3.2
UCB1	5 3.3 PW-UCB1	
	8 3.4 Probabilistic Sharpe Ratio	
	9 3.5 Thompson Sampling	

List of Figures

4.1 Algorithm Comparison	12	4.2 1974-1994
Algorithm Comparison	13	4.3 1994-2020 Algorithm
Comparison	14	

List of Tables

4.1 Evaluation Metrics	11	.1 Yearly
Annualized Sharpe ratios	17	.2 3-Year Moving
Average, Annualized Sharpe ratios	18	

1. Introduction 1

1 Introduction

Historically, there have been many strategies implemented to solve the portfolio choice problem. An accurate estimation of optimal portfolio allocation is challenging due to the non-deterministic complexities of financial markets. Due to this complexity, investors tend to use a mean-variance framework to minimize risk over simply maximizing expected returns. In this paper, we will offer a data-driven approach to allocate wealth to a different set of portfolios such that it maximizes the Sharpe ratio given its context, which in our case context refers to historical market data. One significant difference that is worth pointing out is; Algorithms we experiment in this paper are either non-parametric or take account of the third and fourth moments of return distributions. Meaning, either we do not make any distribution assumption or we take account of the higher moments than first and second. We do not claim the algorithms that are experimented in this paper will yield the same results if they are implemented to live trading environments or when they are applied to different markets or different assets. Nevertheless, our results show that bandit algorithms can learn to act optimally when applied to portfolio choice problem.

Our approach is mostly based on Markowitz (1952) framework. In simple terms, we can summarize our approach as follows; We combine several parametric and non-parametric bandit algorithms with our prior knowledge that we obtain from historical

data. This framework gives us a decision function in which we can choose portfolios to include in our final portfolio. Once we have our candidate portfolio weights, we apply the first-order condition over portfolio variances to distribute our wealth between 2 candidate set of portfolio weights such that it minimizes the variance of the final portfolio.

The rest of the paper is divided as follows: First, we give a brief literature review about methodologies relevant to our approach. Then we provide a detailed description of each algorithm that we experimented with. We point out the advantages and drawbacks of the proposed methods. Finally, we conclude with experimental results and our remarks regarding the result.

2 2. Literature Review

2 Literature Review

If we consider the mean-variance framework, there are several drawbacks to this approach. Using historical returns to estimate the coming period return is not necessarily reliable. Another drawback is that historical variance and covariances do not represent the future behavior of a particular asset or portfolio. But maybe the most strict assumption is normality, it is well known that returns possess a heavy-tailed and skewed distribution, which results in underestimated risk or overestimated returns. In essence, if we can robustly estimate return distribution, we can still use the mean-variance framework to some extent by incorporating higher moments. If we cannot determine the return distribution accurately, we can rely on a non-parametric approach or distribution of different metrics over returns to optimize our portfolio. In this paper, we explore the latter approach.

We start our experiments with Upper Confidence Bound (UCB) algorithm proposed by Auer, Peter (2002). This simple method is proposed to deal with exploration and exploitation problem. This problem can be seen as getting stuck at a local optimum due to not exploring other options. This algorithm gives us a confidence bound for the value expectation. We choose options that maximize the confidence bound. In our framework, these options are different portfolio choices. For the application of UCB to the portfolio choice problem, we follow the framework of Shen (2015). UCB is the non-parametric approach we experiment with. We introduce details in the methodology section. Next, we combine UCB with some results from finance literature. Previously we pointed out that returns do not follow a normal distribution, thus portfolio allocation with mean-variance framework underestimates the risk. Mertens (2002) shows that under IID returns, normality assumption on returns can be dropped, and the Sharpe ratio still follows a normal

distribution. Later Christie (2005) generalized Mertens (2002) result and relaxed the assumption of IID returns. In Christie (2005) derivation, stationarity and ergodicity of returns are sufficient for the normality of the Sharpe ratio. Later on Opdyke (2007) proved that Mertens (2002) and Christie (2005) results are identical. Combining all this work, we end up with the Sharpe ratio that follows a normal distribution and has a closed-form solution for its variance that takes account third and fourth moments of the return distribution. We use these results to incorporate our prior knowledge to UCB. We detail the implementation in the methodology section. In its essence, Bandit algorithms are decision functions that update parameters of their decision function using the data collected from experiments. Nothing is holding us back to experiment with different decision functions outside of the bandit algorithms

2. Literature Review

circle. Like us, Bailey and de Prado (2012) use results on the distribution of the Sharpe ratio, but their approach aims to evaluate hedge fund manager performances. Having this in mind, they derive probabilistic Sharpe ratio that estimates the probability of Sharpe ratio being above a certain threshold. We can pretend that different portfolio allocation strategies are different hedge fund managers and use Bailey and de Prado (2012) results as our decision function, which chooses portfolios. Finally, we use Thompson Sampling, Thompson (1933), which we aim to learn which allocation strategy is the best given their track records; this methodology is also detailed later.

3. Methodology

3 Methodology

We start this section by detailing one of our framework's main components, which are orthogonal portfolios. Obtained portfolio candidates will form the action set of most of the algorithms that we specify in this section.

3.1 Orthogonal Portfolios

We consider daily returns of the 48 value-weighted industry portfolio, which we obtained from Professor French website. We consider gross returns which are given by \mathbb{S}_t where S_t represent the Closing price at time t . Thus our return matrix is given by R , where each column gives the daily return series of each industry portfolio. Considering $k = 48$ different

portfolios for our weight vector, we have:

$$\omega^T \mathbf{1} = 1$$

$$\mathbf{X}_i^k$$

We refer to the covariance matrix of 48 industry portfolios by Σ . Where each Σ estimated from rolling window returns series. Meaning for instance to estimate Σ at time t we use a window from $t - 1$ to $t - 121$. Next, we consider eigenvalue decomposition of the covariance matrix, that is:

$$\Sigma_t = H_t \Lambda_t H_t^T$$

Where H_t is an orthogonal eigenvector matrix and Λ_t is a diagonal matrix with eigenvalues through its diagonal, for time t . One thing to point out here is we sort the eigenvalues in descending order, and we arrange the eigenvector matrix accordingly. Through this decomposition, we obtain k number of uncorrelated portfolios where eigenvalues represent the variances of these portfolios. Furthermore, to meet the constraint that portfolio weights should sum up to 1, we normalize each eigenvector by its L1 norm. New eigenvectors are given

by:

$$\tilde{H}_{t,i} = \frac{H_{t,i}}{H_{t,i}^T \mathbf{1}}$$

Following this, we can calculate the variance of each eigenvector portfolios as

$$\text{follows: } \tilde{\Sigma}_t = \tilde{H}_t \Sigma_t \tilde{H}_t^T = \tilde{\Lambda}_t$$

Since the vectors are orthonormal variance of each of these portfolios is given by normalized

3. Methodology 5

eigenvalues, where return matrix of these portfolios is given by $\tilde{H}_t^T R_t$

Using the expression we obtained above, we can decompose the new covariance matrix as a sum of k rank-one matrices, that is:

$$\tilde{\Sigma}_t = \mathbf{X}_i^l + \sum_{i=l+1}^k \lambda_{t,i} \tilde{H}_{t,i} \tilde{H}_{t,i}^T$$

Where $\lambda_{t,i}$ refers to i^{th} component of sorted eigenvalue matrix at time t . In other words, we followed Principal Component methodology to estimate factors. At its core, we followed the arbitrage pricing theory to decompose the market to its factors. We do not know what

these factors are, but understanding factors is not a necessity for our purposes. So far, we followed the same approach as Shen (2015). From the Arbitrage Pricing Theory, we know looking at the eigenvalues, we can separate factors as market-wide risks and idiosyncratic risks. Some studies consider idiosyncratic risks as pure noise, and it cannot be used to get returns above the market. Mainly we aim to use portfolios that have exposure to idiosyncratic risk to increase Sharpe ratio. In other words, we try to increase our returns without increasing variance too much, such that increase in return compensates the increase in standard deviation. Another question remains on how to determine cutoff point I . One approach would be to follow Fama and French. (1993) results consider 3 or 5 as a cutoff point. We believe factors might be time-varying, and we choose the cutoff point using the following rule:

$$I = 1_{\{median(\lambda) = \lambda_i\}} \cdot i$$

Meaning we look at the location of the median eigenvalue, that we obtained using a rolling window covariance matrix. Next, we introduce algorithms we experimented with and explain their implementations.

3.2 UCB1

Hoeffding's inequality is given by:

$$P(E(X) > \bar{X}_t + u) \leq e^{-2tu^2}$$

When applied to our setting we have:

$$\frac{h}{P} U_t(k_i) \leq e^{-2N_{t(k_i)} U_t(k_i)^2}$$

$$SR(k_i) > \hat{SR}_t + \frac{6}{3. \text{ Methodology}}$$

$SR(k_i)$ refers to the Sharpe ratio of a candidate portfolio, and $U_t(k_i)$ refers to uncertainty around that particular portfolio's Sharpe ratio. As a whole, this equation represents the upper confidence bound of a given portfolio.

Setting this bound to some p and solving for $U_t(k_i)$, we have:

$$\begin{aligned} p &= e^{-2N_{t(k_i)} U_t(k_i)^2} \\ \Rightarrow U_t(k_i) &= \frac{\sqrt{-\log p}}{\sqrt{2N_{t(k_i)}}} \end{aligned} \quad (3.1)$$

Where $N_t(k_i)$ refers to number of times certain orthogonal portfolio chosen. Depending on the problem, upper confidence bound can differ. In our case, we follow the upper confidence bound provided by Shen (2015). Thus we have our policy function: s

$$\pi(k_i) = \frac{\hat{SR}(k_i) + \frac{\tau}{N_t(k_i)}}{2 \log(N_t(k) + \tau)}$$

Where $N_t(k)$ refers to the number of rounds our algorithm traded, τ refers to the size of the training window that we estimate Sharpe ratio and covariance matrix, and $N_t(k_i)$ is the number of times our orthogonal portfolio k_i is selected. Following our policy function, we choose the portfolios with the following rule:

$$k_i^* = \arg \max_i \pi(k_i)$$

Intuitively, we can think of UCB1 as a decision function that tries to maximize its expected return, which is Sharpe ratio. Additionally to that second component of UCB1 takes account of the uncertainty regarding portfolios not chosen before. As algorithm chooses a particular set of portfolio more and more, uncertainty around that particular portfolio decreases. On the other hand if algorithm did not chose particular portfolio often, UCB1 boosts the chance that particular portfolio will be selected through increasing uncertainty around that particular portfolio.

We apply this policy to obtain our two candidate portfolios. The first candidate comes from the set of portfolios that captures the market variation (first / portfolios), and the second candidate comes from the portfolios that represent idiosyncratic risk. From here we apply the first-order condition to distribute our wealth between 2 portfolios such that it minimizes

3. Methodology 7

estimated variance, that is:

$$\lambda_p = \theta_p^2 \tilde{\lambda}_{ki} + (1 - \theta_p)^2 \tilde{\lambda}_{kj}$$

$$\Rightarrow \theta_p^* = \arg \min_{\theta} \lambda_p = \tilde{\lambda}_{kj}$$

$$\tilde{\lambda}_{ki} + \tilde{\lambda}_{kj}$$

$$\omega_p = (1 - \theta_k^*) \tilde{H}_{ki^*} + \theta_k^* \tilde{H}_{kj^*}$$

Steps to implement this algorithm are as follows:

Thus we have our final portfolio weights:

Algorithm 1: UCB1

Inputs: window size τ , return matrix R
(3.2)

Algorithm 1

Initialize: Initialize a reward array and an array that keeps account of how many times a certain portfolio chosen

for t in $\{whole\ sample\ period-\tau\}$ do

 Get the covariance matrix of returns from pre-determined window size; Get the eigenvalue decomposition, sort the eigenvalues and order the eigenvectors accordingly;

 Determine the index of median of eigenvalues l ;

 Normalize the eigenvectors by its L1 norm, then obtain new Covariance matrix;

 Calculate the historical Sharpe ratio of each orthogonal portfolio then apply UCB1 policy function to each portfolio;

 Choose two candidate portfolios by doing argmax over UCB1 policy function, apply first order condition;

 Store the return of chosen portfolio and update the number of times each portfolio has been chosen for selected portfolios in current round;

end

Outputs: The final portfolio weight vector ω_p and the portfolio returns at each time t

8 3. Methodology

3.3 PW-UCB1

One drawback of the UCB1 is, it does not take account of our knowledge about prior distribution, meaning it starts learning how to do portfolio allocation from 0. In most cases this feature of UCB1 is useful due to removing the necessity regarding making a distribution assumption. But if we have some knowledge regarding prior distribution, it may improve our results. We improve on the results of UCB1 by including our knowledge about prior distribution to the algorithm.

We use results from finance literature to include our prior to UCB1's policy function. Consider the result of Mertens (2002), regarding the distribution of Sharpe ratio:

$$\begin{aligned} & \text{"} & + \frac{Y_{4-3}}{4-3} \\ & (\hat{SR}) \sim N & \frac{4SR^2}{n-1} \\ & \# & \\ & \mu_{SR,1} + 0.5 \cdot SR^2 - \gamma_3 SR \end{aligned}$$

Using this distribution we can maximize over:

$$P(SR \leq \hat{SR}_i - SR^*)$$

Assuming Sharpe ratios are coming from the same distribution, we can see that, the greater this probability, the higher the chance a portfolio will achieve a higher Sharpe ratio compared to given threshold. Finally, we weight UCB1 policy function to get our new policy function, later we choose the candidate portfolios according to the following formula:

$$k_i^* = \arg \max_i P(SR \leq \hat{SR}_i - SR^*) \pi(k_i)$$

From here, we again apply First Order Condition and get our final weights. One can ask why we are incorporating UCB1 and instead of maximizing the probability of the Sharpe ratio. The reason is that we are not considering posterior distribution, but rather the unconditional probability distribution of Sharpe ratio. Incorporating UCB1 allows our algorithm to explore different portfolios even if there is another portfolio that has higher $P(SR \geq SR^*)$. For SR^* we use the mean of all candidate portfolio Sharpe ratios, which translates to forcing UCB to form a portfolio at least performs better than average of all candidate portfolios. But this also means taking more risk, which can result in lower returns with specific assets or in times of financial turmoil. This concludes the derivation of the probability-weighted UCB1.

Implementation of this algorithm is same as UCB1, the only difference is to multiply

3. Methodology 9

UCB1 policy function by probabilities obtained from above distribution using CDF function of standard normal distribution.

3.4 Probabilistic Sharpe Ratio

As we mentioned before, Bailey and de Prado (2012) use the same results to derive a metric called probabilistic Sharpe ratio to evaluate hedge fund manager performances, we again select two candidate portfolios and use the first-order condition over that portfolio,

but this time our policy function is probabilistic Sharpe ratio proposed by Bailey and de Prado (2012). Probabilistic Sharpe ratio is given by:

$$P(SR \leq SR^*) = \frac{\Phi\left(\frac{SR^* - \hat{SR}}{\sqrt{\frac{1}{n-1} \sum_{t=1}^n (SR_t - \hat{SR})^2}}\right)}{\Phi\left(\frac{SR^* - \hat{SR}}{\sqrt{\frac{1}{n-1} \sum_{t=1}^n (SR_t - \hat{SR})^2}}\right) + \frac{1}{4} SR^2}$$

We again take threshold SR^* as the mean of all candidate portfolio Sharpe ratios. Implementation is again same, the only difference is to replace UCB1 policy function with probabilistic Sharpe ratio.

3.5 Thompson Sampling

Thompson sampling uses the Beta-Bernoulli model to select actions. In our context, we use Thompson sampling to choose different portfolio strategies. Candidate strategies given to Thompson Sampling approach are the algorithms introduced so far plus additional strategies, namely minimum variance portfolio and equal weight portfolio. We consider each strategy outcome at certain t as a success or failure, which results in Bernoulli distribution, from there, we know Beta distribution is the conjugate distribution to Bernoulli that is:

$$P(X = 1|D) \sim \text{Beta}(\alpha, \beta)$$

Where each round we update parameters of beta distribution(α and β), if certain candidate strategy achieves maximum return among all candidate strategies in a particular period, that round is counted as a success for that strategy and new α is given by: $\alpha_{k,t} = \alpha_{k,t-1} + 1$ whereas for other strategies are counted as failures which results in increasing the β .

Implementation is as follows:

10 3. Methodology

Algorithm 2

Algorithm 2: Thompson Sampling

Inputs: window size τ , return matrix R

Initialize: Initialize a reward array, the array that keeps account of how many times a certain portfolio has been chosen, and 2 additional arrays that keeps track of the number of times a strategy succeed or failed respectively

for t in $\{whole\ sample\ period - \tau\}$ do

Follow all the steps in above algorithms to obtain 3-different set of portfolio weights, also include another weights set that gives equal weight to all industry portfolios; Draw probabilities for each candidate strategy from beta distribution which alpha and beta parameters are the number of times a strategy succeed or failed;
Choose your final strategy by; $\arg \max_i P(X_i = 1 | Data)$, where $X_i = 1$ refers to success;
Store the return of chosen portfolio, update all the arrays;

end

Outputs: The final portfolio weight vector ω_p and the portfolio returns at each time t

With this, we conclude the introduction of all algorithms experimented in this paper. Implementation of all algorithms conducted in Python, one can find code and results presented in this paper at authors' github page. Next, we move on to experimental results.

4. Results

4 Results

Our experiments are based on 48 US value-weighted industry portfolios. We consider the time range 1974-02 to 2019-12; we avoid 2020 due to Covid-19. The below table presents the metrics we used to evaluate strategy performances. Below the table, one can observe the evaluation of cumulative wealth through the whole investment period. In the appendix, we provide mean annualized Sharpe ratios for every year in our time range. As metrics, we use cumulative wealth, the wealth accumulated through years by the strategy assuming starting from 1 Dollar. We also report the annualized Sharpe ratios in the whole period and mean of yearly annualized Sharpe ratios. Below table reports the results of following strategies by order; Minimum Variance Portfolio (MVP), Constant Weight Rebalance portfolio (CWR), Equal Weight portfolio (EW), Upper Confidence Bound 1 (UCB1), Thompson Sampling (TS), Maximum Probabilistic Sharpe ratio (MaxPSR), Probability Weighted UCB1 (PW-UCB1).

Table 4.1: Evaluation Metrics


Descriptive Statistics

	Returns(%)				Dev.(%)				Annualize				Wealt			
	Annualize				Annualize				d Mean				h			
	d				d Mean				SR							
	Mea				St				Cumulativ							
	n				d				e							
					Yearly											
MVP	8.46	11.22	0.75	0.99	35.32	CWR	13.06	16.17	0.81	1.05	208.6	EW	12.87	15.81	0.81	1.04
	12.87	15.81	0.81	1.04	196.53	UCB1	16.31	17.83	0.91	1.07	818.17	TS	13.67	16.9	0.81	0.96
	13.67	16.9	0.81	0.96	262.53	MaxPSR	11.93	21.84	0.55	0.74	78.23					

One thing to observe here, even UCB1 and PW-UCB, yield the highest Sharpe ratios. They also have the highest standard deviation, which implies bandit portfolios tend to take more risk than methodologies that aim to minimize variance, but this was already expected due to the exploration component. Our purpose was to see if the bandit strategy can increase the return such that it offsets the increase in standard deviation. Thompson sampling yields a lower standard deviation because TS also consists of portfolio strategies that aim to minimize variance in its action set.

12 4. Results

Figure 4.1: Algorithm Comparison




In this part, we include the evaluation of cumulative wealth throughout the whole period and in 10 year time intervals. One interesting thing to notice is that bandit algorithms' performance diminishes during periods of high momentum followed by turmoil. The drop in the bandit algorithms' cumulative wealth is more severe compared to classic allocation strategies such as *EW* or *MV P*. Especially PW-UCB1, this also can be seen from standard deviation of the returns. This is due to using the rolling window to estimate moments of the return distribution. Since we are weighing UCB1 with the Sharpe ratio probability, and since this probability reflects the 120-day window, algorithm puts more weights on industries that gain more during high momentum periods, such as technology portfolio. During the dot.com bubble(1995-2002) period, UCB1 and PW-UCB1 gain a lot by putting more weight on

technology portfolio, but they suffer the most, during the turmoil that followed high momentum period.


4. Results 13

Figure 4.2: 1974-1994 Algorithm Comparison



14 4. Results

Figure 4.3: 1994-2020 Algorithm Comparison



5. Conclusion 15

5 Conclusion

In this project, we have studied the multi-armed bandit problem as a mathematical model for sequential decision making under uncertainty. In particular, we focus on its application in financial markets and construct a high-quality portfolio selection algorithm.

Considering that we did not take into account transaction costs and did only backtesting rather than live trading environment, one cannot use the algorithm for real-time trading purposes without further improvements. Nevertheless, the results of experiments on FF48 showed that portfolio selection strategy based on our algorithm outperforms classic strategies such as EW or MVP in terms of annualized mean returns, Sharpe ratio and

cumulative wealth. The reason is that our algorithm allows dynamic asset allocation with the relaxation of strict normality assumption on returns and incorporates Sharpe ratio probability to better evaluate performances. Our algorithm may suffer from crashes during the turmoils by putting more weight on assets whose past performances are well due rolling window estimation of the covariance matrix and returns vector, one can solve this issue by using more sophisticated prediction model to estimate returns and the covariance matrix. To conclude, our algorithm could appropriately balance the benefits and risks well and achieve higher returns by controlling risk when the market is stable.

16 References **References**

Auer, Peter (2002). Using confidence bounds for exploitation-exploration trade-offs. Bailey, D. H. and de Prado, M. L. (2012). The sharpe ratio efficient frontier. Christie, S. (2005). Is the sharpe ratio useful in asset allocation? Fama, E. F. and French., K. R. (1993). Common risk factors in the returns on stocks and bonds. Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1):77–91. Mertens, E. (2002). Comments on variance of the iid estimator in lo. Opdyke, J. D. J. (2007). Comparing sharpe ratios: so where are the p-values? Shen, W. (2015). Portfolio choices with orthogonal bandit learning. Thompson, W. R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples.

Appendices

Table .1: Yearly Annualized Sharpe ratios

	MVP	CWR	EW	UCB1	TS	MaxPSR	PW-UCB1
1975	0.74	0.56	0.56	0.64	0.25	0.28	0.58
1976	3.09	2.47	2.47	2.17	1.98	1.84	2.21
1977	1.81	1.4	1.49	1.8	1.78	0.92	1.79
1978	-1.22	0.61	0.71	-0.22	-0.65	0.02	-0.2
1979	-1.44	1.6	1.67	-0.42	-0.06	0.17	-0.48
1980	0.35	2.06	2.17	0.93	1.42	1.07	1.02
1981	1.79	1.08	1.08	1.3	0.91	1.36	1.21
1982	3.09	0.86	0.82	1.57	0.41	0.99	1.85
1983	3.55	1.59	1.53	2.37	0.79	1.78	2.83

1984	2.71	0.93	0.85	1.95	1.12	1.7	1.9
1985	2.96	1.29	1.26	2.1	2.15	1.61	2.1
1986	2.04	1.94	1.73	2.21	2.35	1.61	2.07
1987	-0.08	0.52	0.38	1.09	1.3	0.49	1.33
1988	-0.03	0.58	0.52	0.74	0.73	0.9	1.0
1989	1.37	1.68	1.81	0.94	0.87	1.48	0.77
1990	-0.47	0.58	0.77	1.29	0.58	-0.33	1.11
1991	-0.18	0.83	0.96	1.28	1.25	-0.07	1.16
1992	1.62	1.85	1.76	1.33	2.39	1.03	1.34
1993	1.61	1.34	1.02	0.99	1.5	1.17	1.08
1994	0.9	0.68	0.56	0.49	0.75	0.95	0.48
1995	1.63	1.62	1.64	1.25	1.36	1.36	1.25
1996	3.14	2.77	2.63	2.6	2.67	2.77	2.63
1997	2.17	1.76	1.66	3.02	2.79	1.98	3.16
1998	0.42	1.04	1.05	1.88	1.36	0.66	1.86
1999	-1.25	0.55	0.6	0.81	0.71	-0.21	0.42
2000	-1.67	0.56	0.4	-0.05	0.34	-0.46	-0.18
2001	-0.62	0.39	0.07	-0.66	-0.6	-0.28	-0.57
2002	0.04	-0.13	-0.35	0.11	-0.26	-0.27	0.29
2003	0.67	0.62	0.45	0.59	0.54	0.53	0.85
2004	2.05	1.78	1.57	1.19	1.54	1.52	1.34
2005	0.8	1.02	1.01	0.82	1.11	0.67	1.06
2006	0.94	1.01	1.12	0.65	0.65	0.51	0.82
2007	1.09	0.98	1.02	0.35	0.09	0.44	0.54
2008	-1.1	-0.44	-0.5	-0.15	-0.22	-1.17	-0.08
2009	-0.84	-0.02	-0.15	0.01	0.3	-0.39	-0.02
2010	0.56	1.18	1.14	0.59	1.26	1.73	0.77
2011	1.14	0.54	0.62	1.19	1.26	1.15	1.53
2012	0.76	0.46	0.55	0.54	0.46	0.27	0.43
2013	1.76	1.78	1.97	1.0	1.05	1.1	0.61
2014	3.02	1.56	1.85	1.63	0.96	0.81	1.5
2015	0.93	0.11	0.54	0.51	-0.11	-0.3	0.49
2016	0.93	0.44	0.6	1.57	1.0	-0.17	1.34
2017	2.32	1.65	1.66	2.56	2.07	0.55	2.35
2018	0.54	0.4	0.43	0.56	0.31	0.66	0.65

Table .2: 3-Year Moving Average, Annualized Sharpe ratios

MVP CWR EW UCB1 TS MaxPSR PW-UCB1

1979	0.6	1.33	1.38	0.79	0.66	0.65	0.78
1980	0.52	1.63	1.7	0.85	0.89	0.81	0.87
1981	0.26	1.35	1.42	0.68	0.68	0.71	0.67
1982	0.51	1.24	1.29	0.63	0.4	0.72	0.68
1983	1.47	1.44	1.46	1.15	0.69	1.07	1.28
1984	2.3	1.31	1.29	1.62	0.93	1.38	1.76
1985	2.82	1.15	1.11	1.86	1.08	1.49	1.98
1986	2.87	1.32	1.24	2.04	1.36	1.54	2.15
1987	2.23	1.25	1.15	1.94	1.54	1.44	2.05
1988	1.52	1.05	0.95	1.62	1.53	1.26	1.68

1989 1.25 1.2 1.14 1.42 1.48 1.22 1.45
1990 0.57 1.06 1.04 1.25 1.16 0.83 1.26
1991 0.12 0.84 0.89 1.07 0.95 0.49 1.07
1992 0.46 1.1 1.16 1.12 1.17 0.6 1.08
1993 0.79 1.26 1.26 1.17 1.32 0.65 1.09
1994 0.7 1.06 1.01 1.08 1.3 0.55 1.03
1995 1.12 1.26 1.19 1.07 1.45 0.89 1.06
1996 1.78 1.65 1.52 1.33 1.73 1.45 1.36
1997 1.89 1.63 1.5 1.67 1.81 1.64 1.72
1998 1.65 1.58 1.51 1.85 1.79 1.54 1.88
1999 1.22 1.55 1.52 1.91 1.78 1.31 1.87
2000 0.56 1.34 1.27 1.65 1.57 0.95 1.58
2001 -0.19 0.86 0.76 1.0 0.92 0.34 0.94
2002 -0.62 0.48 0.36 0.42 0.31 -0.11 0.36
2003 -0.57 0.4 0.24 0.16 0.15 -0.14 0.16
2004 0.09 0.64 0.43 0.24 0.31 0.21 0.35
2005 0.59 0.74 0.55 0.41 0.47 0.43 0.59
2006 0.9 0.86 0.76 0.67 0.72 0.59 0.87
2007 1.11 1.08 1.03 0.72 0.79 0.73 0.92
2008 0.76 0.87 0.84 0.57 0.63 0.39 0.73
2009 0.18 0.51 0.5 0.34 0.39 0.01 0.46
2010 0.13 0.54 0.53 0.29 0.42 0.22 0.41
2011 0.17 0.45 0.42 0.4 0.54 0.35 0.55
2012 0.1 0.34 0.33 0.44 0.61 0.32 0.53
2013 0.68 0.79 0.82 0.67 0.87 0.77 0.66
2014 1.45 1.1 1.22 0.99 1.0 1.01 0.97
2015 1.52 0.89 1.1 0.97 0.72 0.6 0.91
2016 1.48 0.87 1.1 1.05 0.67 0.34 0.88
2017 1.79 1.11 1.32 1.45 0.99 0.4 1.26
2018 1.55 0.83 1.01 1.36 0.84 0.31 1.27