The Kelly Growth Optimal Portfolio with Ensemble Learning*

Weiwei Shen,[1,2] Bin Wang,[3] Jian Pu,[1] Jun Wang[1]
[1]Shanghai Key Lab for Trustworthy Computing, School of Computer Science and Software Engineering
East China Normal University, Shanghai, China
[2]GE Global Research Center, Niskayuna, NY, USA
[3]Credit and Portfolio Risk Management, Citigroup, St. Louis, MO, USA
realsww@gmail.com, vince.wang@citi.com, jianpu@sei.ecnu.edu.cn, jwang@sei.ecnu.edu.cn

Abstract

As a competitive alternative to the Markowitz mean-variance portfolio, the Kelly growth optimal portfolio has drawn sufficient attention in investment science. While the growth optimal portfolio is theoretically guaranteed to dominate any other portfolio with probability 1 in the long run, it practically tends to be highly risky in the short term. Moreover, empirical analysis and performance enhancement studies under practical settings are surprisingly short. In particular, how to handle the challenging but realistic condition with insufficient training data has barely been investigated. In order to fill voids, especially grappling with the difficulty from small samples, in this paper, we propose a growth optimal portfolio strategy equipped with ensemble learning. We synergically leverage the bootstrap aggregating algorithm and the random subspace method into portfolio construction to mitigate estimation error. We analyze the behavior and hyperparameter selection of the proposed strategy by simulation, and then corroborate its effectiveness by comparing its out-of-sample performance with those of 10 competing strategies on four datasets. Experimental results lucidly confirm that the new strategy has superiority in extensive evaluation criteria.

1 Introduction

How to wisely construct an investment portfolio with access to multiple assets has been a long-lasting question of great concern to both individual and institutional investors (Brandt 2010). The seminal work by (Markowitz 1952) presents a mean-variance framework as one prevalent answer to this question. Meanwhile, the growth optimal portfolio, which is originally considered for gambling by (Kelly 1956) and introduced into investment by (Latané 1959), offers a valuable alternative. In addition, various successful investors, such as Warren Buffet, James Simons and John Maynard Keynes, are reported to have adopted Kelly-type strategies to manage their funds (MacLean, Thorp, and Ziemba 2011). Briefly, different from the myopic and risk-return tradeoff consideration in the mean-variance framework, the Kelly growth optimal portfolio is designed to maximize the long-term exponential growth rate of an investment capital over multiple periods. On the one hand, in mild conditions it has been theoretically proved dominating the value of any other portfolio over the long haul and requiring the least time to reach a preassigned wealth target (Breiman 1961). On the other hand, however, with the advent of the growth optimal portfolio its overly high risk in the short term serves as one major source of enduring debates on its merit as a normative investment rule in portfolio research (Samuelson 1979; Ziemba 2015). Moreover, in contrast to the unremitting efforts expended to improve the poor out-of-sample performance for the classical mean-variance portfolio (Kolm, Tütüncü, and Fabozzi 2014), empirical studies and enhancing methods for the growth optimal portfolio are surprisingly scarce (Estrada 2010; Bottazzi and Santi 2017).

In portfolio construction, balancing stationarity and estimation accuracy of model input parameters has been one practical conundrum (Broadie 1993). On the one hand, since return data from 20 years ago might have little bearing on returns this year, parameters are unlikely to be stationary over a long period of time. On the other hand, while the estimation errors decrease when more data are used, more data demand a longer time horizon when the data frequency is fixed. To overcome this dilemma between stationarity and accuracy, one oft-quoted approach in Markowitz's framework is the resampled mean-variance portfolio proposed in (Michaud 1989). In effect, it applies a parametric bootstrap aggregating algorithm to reduce estimation risk and advance out-of-sample performance. However, its efficacy has been continually called into question due to its limited or even no improvement over Markowitz's portfolio (Wolf 2013).

As yet, to the best of our knowledge, no work exists on applying appropriate ensemble methods into Kelly's portfolio to boost its performance. Intuitively, as the growth optimal strategy essentially relies on the complete knowledge of the joint distribution of asset returns, insufficient data for estimation would dramatically deteriorate its performance. Its degeneration is more severe than that of the mean-variance portfolio, where only the first two moments of asset returns need estimation. Meanwhile, ensemble learning as a rich area for tools has achieved impressive success in heightening performance of a vast class of existing algorithms (Zhou 2012). Hence, how to exploit ensemble learning in Kelly's portfolio deserves a thorough and fresh investigation.

Motivated by the preceding consideration, in this paper, we propose an ensemble growth optimal portfolio strategy.

*This work was supported in part by NSFC (No. 61672236).

Specifically, we concentrate on facilitating the applicability of the Kelly growth optimal portfolio from a practical standpoint when the number of assets and the number of historical asset return data are close. In our study, we synergically combine the bootstrap aggregating algorithm (Breiman 1996) and the random subspace method (Ho 1998) in the Kelly growth optimal portfolio to mitigate estimation risk rooted in a small sample size. In particular, the new strategy is composed of two loops. The outer loop follows a parametric bootstrap aggregating algorithm. In each iteration it calls the inner loop to generate one basis portfolio and then takes the average over all the generated basis portfolios to determine the final portfolio strategy. The inner loop employs a subset resampling method illuminated by the random subspace method. For each desired basis portfolio it averages over multiple estimated optimal weights from resampled small portfolios. Such an approach targets at reducing estimation errors by sacrificing some diversification benefits. For validation we analyze its behavior and hyperparameter selection in a simulation study, and then empirically compare its out-of-sample performance with those of 10 other strategies on four datasets. Experimental results evidently show that the new strategy has considerable superiority in extensive evaluation criteria.

2 Background and Related Work

In this section, we review finance and machine learning papers on the Kelly growth optimal portfolio, and then recapitulate relevant observations in the Markowitz mean-variance portfolio and ensemble methods that motivate our work.

The growth optimal portfolio, also called the log-optimal portfolio or Kelly's criterion, makes investment decisions by maximizing the growth rate of the invested capital (Kelly 1956; Latané 1959). It is equivalent to maximizing the geometric mean return or the median wealth of a portfolio (Luenberger 1998). However, limited work has been undertaken from empirical and performance enhancing aspects. Specifically, (Hunt 2005) shows that the growth optimal portfolio is commonly unattainable due to huge leveraging and volatility. (Estrada 2010) concludes that the growth optimal portfolio is exceedingly aggressive for short-term investors. (Bottazzi and Santi 2017) demonstrate that among 11 strategies including the growth optimal portfolio no clear winner emerges. A review may be referred to (Christensen 2005).

Meanwhile, the successful records of leveraging machine learning algorithms into numerous regimes have promoted their roles in portfolio research as a new trend (Kolm, Tütüncü, and Fabozzi 2014). Over years machine learning researchers have been deeply inspired by Kelly's criterion in portfolio construction. Among them, (Cover 1991) in his seminal work proposes a universal portfolio strategy based on perfect knowledge of the future without making any assumption on the assets return dynamics. It is analogous to Kelly's portfolio in hindsight. (Kalai and Vempala 2002) present an efficient implementation of Cover's algorithm to retard the exponential growth of computational costs with the number of assets for universal portfolios. (Agarwal et al. 2006) consider an online Newton step algorithm to compute the portfolio in the universal portfolio context. (Györfi, Ottucsák, and Walk 2012) publish a self-contained text contributing its main content on Kelly's portfolio with machine learning. (Shen et al. 2015) and (Shen and Wang 2016) introduce bandit learning and Thompson sampling into portfolio construction, respectively. (Li et al. 2017) provide an online learning portfolio framework in the presence of transaction costs. (Uziel and El-Yaniv 2018) underscore its asymptotic properties with tail risk constraints under unknown stationary and ergodic processes. An in-depth review may be found in the monograph by (Li and Hoi 2015).

On the one hand, the well-known mean-variance portfolio often has poor out-of-sample performance, especially when the number of assets is larger than that of data (Broadie 1993). That is primarily because estimation errors in input parameters, i.e., the first two moments of asset returns, are amplified by optimization and then propagate into the solution (Michaud 1989). Among efforts in retrofitting its performance, a strand of work respectively design two-fund, three-fund and four-fund aggregating portfolios to cancel propagated errors in optimal solutions by a handful of sophisticated basis portfolios (Kan and Zhou 2007; Tu and Zhou 2011; Kan, Wang, and Zhou 2016). Their experiments on moderately large samples have substantiated the error cancellation idea through boosted performance. On the other hand, ensemble learning virtually shares the theme of diversification with portfolio research (Derbeko, El-Yaniv, and Meir 2002). It replaces financial assets by base learners from algorithms. The level of diversity of individual learners determines the generalization quality of the aggregated learner. In particular, (Shen and Wang 2017) apply a subset resampling algorithm into the mean-variance portfolio and obtain promising results. Thus, we believe that ensemble learning could be a more general and effective approach for canceling estimation errors in Kelly's portfolio than aforementioned multiple fund aggregations .

Further, among ensemble methods, the bootstrap aggregating algorithm (bagging) and the random subspace method are particularly efficacious in improving weak learners when the training set is small (Rokach 2010; Polikar 2006). For a decreasing learning curve, i.e., that the generalization error of the base learner decreases with an increase in the training sample size, the random subspace method is commonly recommended. For a flat and non-decreasing learning curve, bagging could be advantageous when the sample size is critical (Skurichina and Duin 2002). As Kelly's portfolio implicitly requires estimating a high-dimensional distribution of asset returns, estimation errors in all the moments and comoments of the distribution will be amplified by optimization and then affect the solution. If the distribution is statistically stationary, implying that the learning curve would be decreasing, then the random subspace method will take effect. In addition, if the decreasing rate of the learning curve is low ascribed to regime shift or high volatility reflected in the selected data span, then bagging may be instrumental in the out-of-sample performance as well. Noticeably, (Kleiner et al. 2014) present a scalable bootstrap for massive data by combining multiple ensemble methods. Therefore, for improving the Kelly growth optimal portfolio, we attempt to infuse bagging and the random subspace method.

3 Methodology

In this section, we first introduce notations and finance terms. Then we describe the proposed portfolio selection strategy. Finally we discuss properties and behaviors of the new method in a simulation study.

Notations

In a frictionless, self-financing, discrete-time and finite horizon investment environment, we denote a series of trading periods as $t_k = k\Delta t$, $k = 0, \ldots, m$, where Δt represents one week or one month, depending upon the rebalancing interval in the context. For simplicity, we use k as the index to indicate the trading period at time t_k hereinafter. From time t_{k-1} to t_k the gross return vector of n risky assets accessible to investors is denoted as $\mathbf{R}_k = (R_{k,1}, \ldots, R_{k,n})^\top$. The gross return $R_{k,i}$ for the i-th asset is computed as $R_{k,i} = S_{k,i}/S_{k-1,i}$, where $S_{k,i}$ and $S_{k-1,i}$ represent the prices of the i-th asset at time t_k and t_{k-1}, respectively. Denote by $\boldsymbol{\omega}_k = (\omega_{k,1}, \ldots, \omega_{k,n})^\top$ the vector of the portfolio weights as the investment decision at time t_k. The i-th element of $\boldsymbol{\omega}_k$ specifies the invested percentage of wealth in the i-th asset. The completion investment condition reads $\boldsymbol{\omega}_k^\top \mathbf{1} = 1$, where $\mathbf{1}$ stands for the $n \times 1$ vector of ones. $\omega_{k,i} > 0$ means that investors take a long position of the i-th asset; $\omega_{k,i} < 0$ indicates a short sale of the i-th asset. For a short sale position, investors sell the borrowed i-th asset for cash and then invest it in other assets. If the price of the borrowed asset surges, investors who have obligation to buy back and return the borrowed asset will suffer from a loss. The realized portfolio net return r_k from time t_{k-1} to t_k is computed as $r_k = \mathbf{R}_k^\top \boldsymbol{\omega}_{k-1} - 1$.

Ensemble Growth Optimal Portfolio

At time t_k investors construct the unconstrained growth optimal portfolio based on the following optimization:

$$\max_{\boldsymbol{\omega}_k} \mathbb{E}_k[\ln(\boldsymbol{\omega}_k^\top \mathbf{R}_k)] \quad \text{s.t.} \quad \boldsymbol{\omega}_k^\top \mathbf{1} = 1, \tag{1}$$

where the symbol $\mathbb{E}_k[\cdot]$ denotes the conditional expectation up to time t_k. As portfolio position constraints are generally equivalent to some shrinkage estimators that can unintentionally curtail estimation risk (Jagannathan and Ma 2003), no position constraints are imposed for obtaining an uncontaminated assessment of the proposed method.

In order to compute the conditional expectation, we need to know the multivariate probability distribution of asset returns. Practically, given τ number of periods of return data up to time t_k as $\{\mathbf{R}_l\}_{l=k-\tau+1}^{k}$, assuming the returns are identically and independently distributed (i.i.d.), the optimization (1) can be approximated by

$$\max_{\boldsymbol{\omega}_k} \sum_{l=k-\tau+1}^{k} \frac{1}{\tau} \ln(\boldsymbol{\omega}_k^\top \mathbf{R}_l) \quad \text{s.t.} \quad \boldsymbol{\omega}_k^\top \mathbf{1} = 1, \tag{2}$$

where the continuous distribution of asset returns is discretely approximated by a probability mass function (Uziel and El-Yaniv 2018; Bottazzi and Santi 2017). The objective function in (2) with a logarithmic function implicitly rests on sample estimates of all the moments and comoments of the distribution from the data $\{\mathbf{R}_l\}_{l=k-\tau+1}^{k}$. While approximating a high-dimensional distribution by a small set of sample, i.e., $\tau \approx n$, is formidably challenging, the accurate quantification of comovements in asset returns is critical in portfolio construction. As such, in our study, we apply ensemble learning to lessen the impact of the associated estimation errors on the optimal solution.

Algorithm 1 Ensemble Growth Optimal Portfolio

Inputs: τ: number of periods of return data; n: number of assets; $\{\mathbf{R}_l\}_{l=k-\tau+1}^{k}$: historical return data for training; $\mathbf{R}_{k+1}$: one out-of-sample return for testing; n_1: number of resamples; n_2: size of each resample; n_3: number of resampled subsets; n_4: size of each subset;
Compute the sample covariance matrix $\hat{\boldsymbol{\Sigma}}_k$ and the sample mean $\hat{\boldsymbol{\mu}}_k$ of the return data $\{\mathbf{R}_l\}_{l=k-\tau+1}^{k}$;
for $h = 1 \to n_1$ **do**
  Generate n_2 normally distributed returns of all the n assets $\{\tilde{\mathbf{R}}_l^h\}_{l=1}^{n_2}$ based on the parameters $\hat{\boldsymbol{\Sigma}}_k$ and $\hat{\boldsymbol{\mu}}_k$;
  for $j = 1 \to n_3$ **do**
    Randomly sample a set $\mathcal{I}_j$ of n_4 indices from $\{1, \ldots, n\}$ without replacement;
    Select the return data $\{\tilde{\mathbf{R}}_l^{h,j}\}_{l=1}^{n_2}$ from $\{\tilde{\mathbf{R}}_l^h\}_{l=1}^{n_2}$ according to the index set $\mathcal{I}_j$;
    Compute the optimal subset portfolio weights $\tilde{\boldsymbol{\omega}}_k^{h,j}$ by solving optimization (2) with $\{\tilde{\mathbf{R}}_l^{h,j}\}_{l=1}^{n_2}$;
    Construct the weights for the basis portfolio $\hat{\boldsymbol{\omega}}_k^{h,j} = (\hat{\omega}_{k,1}^{h,j}, \ldots, \hat{\omega}_{k,n}^{h,j})^\top$ through
    for $i = 1 \to n$ **do**
      $\hat{\omega}_{k,i}^{h,j} = \tilde{\omega}_{k,i}^{h,j} \mathbb{I}\{i \in \mathcal{I}_j\}$;
  Construct the basis portfolio from the n_3 subsets as $\hat{\boldsymbol{\omega}}_k^h = n_3^{-1} \sum_{j=1}^{n_3} \hat{\boldsymbol{\omega}}_k^{h,j}$;
Aggregate the basis portfolios as $\hat{\boldsymbol{\omega}}_k = n_1^{-1} \sum_{h=1}^{n_1} \hat{\boldsymbol{\omega}}_k^h$;
Compute the realized out-of-sample portfolio net return $\hat{r}_{k+1} = \mathbf{R}_{k+1}^\top \hat{\boldsymbol{\omega}}_k - 1$;
Outputs: The vector of portfolio weights $\hat{\boldsymbol{\omega}}_k$ and the realized out-of-sample portfolio return $\hat{r}_{k+1}$.

The proposed ensemble growth optimal portfolio (EGO) presented in Algorithm 1 is straightforward to implement. The algorithm consists of an outer loop and an inner loop. At each time t_k, the outer loop of EGO follows a parametric bagging algorithm that repeats the basis portfolio construction in the inner loop for n_1 times and then takes the average over the generated n_1 basis portfolios to produce the final portfolio strategy. The inner loop applies a subset resampling approach inspired from the random subspace method to construct each basis portfolio by averaging over n_3 optimal weights from n_3 resampled small portfolios.

Specifically, in the h-th outer iteration for $h = 1, \ldots, n_1$, EGO first generates n_2 normally distributed returns of all n assets $\{\tilde{\mathbf{R}}_l^h\}_{l=1}^{n_2}$ based on the sample covariance matrix $\hat{\boldsymbol{\Sigma}}_k$ and the sample mean of returns $\hat{\boldsymbol{\mu}}_k$ estimated from the historical data $\{\mathbf{R}_l\}_{l=k-\tau+1}^{k}$. Second, using $\{\tilde{\mathbf{R}}_l^h\}_{l=1}^{n_2}$ as the input sample, the inner loop of EGO as the basis portfolio constructor produces the h-th basis portfolio $\hat{\boldsymbol{\omega}}_k^h$. Third, EGO attains the portfolio strategy by averaging over the derived n_1 basis portfolios $\hat{\boldsymbol{\omega}}_k = n_1^{-1} \sum_{h=1}^{n_1} \hat{\boldsymbol{\omega}}_k^h$. Accordingly, the realized out-of-sample portfolio net return $\hat{r}_{k+1}$ from time t_k to t_{k+1} is computed as $\hat{r}_{k+1} = \mathbf{R}_{k+1}^\top \hat{\boldsymbol{\omega}}_k - 1$.

In the j-th inner iteration for $j = 1, \ldots, n_3$, EGO uniformly at random samples n_3 subsets of size $n_4 \leq n$ from the original n assets. Denote by $\mathcal{I}_j \subset \{1, \ldots, n\}$ the corresponding index set with $|\mathcal{I}_j| = n_4$. Denote by $\{\tilde{\mathbf{R}}_l^{h,j}\}_{l=1}^{n_2}$ the associated return data according to the index set $\mathcal{I}_j$. The selected subset of returns $\{\tilde{\mathbf{R}}_l^{h,j}\}_{l=1}^{n_2}$ has a size of $n_2 \times n_4$ with the typical size order of $n_4 \leq n \leq n_2$. Next, for each subset, EGO computes the optimal subset portfolio weights $\tilde{\boldsymbol{\omega}}_k^{h,j}$ by solving the optimization problem (2) based on the lower-dimensional data $\{\tilde{\mathbf{R}}_l^{h,j}\}_{l=1}^{n_2}$. This step essentially decreases the dimensionality of the distribution involved in optimization from n to n_4. Then, the inner loop averages over portfolio weights from all the subsets to generate the weights for the basis portfolio as $\hat{\boldsymbol{\omega}}_k^h = n_3^{-1} \sum_{j=1}^{n_3} \hat{\boldsymbol{\omega}}_k^{h,j}$, where $\hat{\boldsymbol{\omega}}_k^{h,j} = (\hat{\omega}_{k,1}^{h,j}, \ldots, \hat{\omega}_{k,n}^{h,j})^\top$ is a prolonged $n \times 1$ column vector with its element as $\hat{\omega}_{k,i}^{h,j} = \tilde{\omega}_{k,i}^{h,j} \mathbb{I}\{i \in \mathcal{I}_j\}$, and the symbol $\mathbb{I}(\cdot)$ stands for an indicator function.

In particular, if $n_1 = 1$ and $n_2 = \tau$, the outer loop degenerates and only the subset resampling step in the inner loop remains functional. We term this case as the subset resampled growth optimal portfolio (SS) for simplicity. If $n_3 = 1$ and $n_4 = n$, the inner loop is redundant and only bagging in the outer loop remains effective. We term this case as the resampled growth optimal portfolio (RE) for brevity. Likewise, if $n_1 = 1$, $n_2 = \tau$, $n_3 = 1$ and $n_4 = n$, EGO degenerates to the classical growth optimal portfolio (GO).

Discussion

EGO collaboratively integrates bagging in the outer loop and subset resampling in the inner loop to alleviate the estimation risk from the parameter uncertainty by sacrificing some diversification benefits. The pivotal hyperparameters that determine its performance characteristics include the number of resamples n_1, the size of each resample n_2, the number of resampled subsets n_3 and the size of each subset n_4. The embedded structure of the two loops assists with their selection. The former two are demanded in the outer loop; the latter two are needed in the inner loop. On the one hand, if available computing power is unrestricted, sufficiently large values of n_1, n_2 and n_3 should be preferred to achieve convergence. Thus, sensitivity analysis for those three hyperparameters will be emphasized. On the other hand, however, a tradeoff exists when we determine n_4. The smaller the subset size is, the more accurate the estimation of the n_4-dimensional joint distribution would be, yet the more diversification benefits within the basis portfolio would get lost, and vice versa. Undoubtedly, a large value of n_3 increases diversification across basis portfolios. Thus, while general results for the optimal hyperparameters are unavailable as they would hinge on the underlying return dynamics of assets, our study below offers some insights and suggestions.

In Figure 1, we investigate the empirical behavior of EGO with respect to the four hyperparameters in constructing growth frontiers by simulation. Briefly, a growth frontier plots the best growth and risk tradeoff curve that a set of assets could possibly achieve by following one given strategy (Luenberger 1998). One growth frontier dominates another when the former lies on the upper left of the latter. Following (Broadie 1993), we report two types of frontiers: actual and true frontiers. An actual frontier represents out-of-sample performance, whereas a true frontier stands for the best possible performance. In the case with a lognormal return model the growth optimal portfolio has a closed-form solution, so we can construct the true frontiers in our simulation study.

From Figure 1a to Figure 1l, we analyze the two special cases, i.e., RE and SS, to shed light on the more complicated case of EGO discussed afterwards. First, those figures indicate that the ensemble step in the inner loop can improve performance more than that in the outer loop. Compared with the growth frontiers generated by RE, those by SS locate closer to the true frontiers, and always dominate the frontiers generated by GO. Second, Figures 1a, 1e, 1i, 1b, 1f and 1j illustrate that while repeating the outer loop for a small number of times, e.g., $n_1 = 50$, may be sufficient, the length of each resample should be longer than that of the input data, e.g., $n_2 = 3\tau$. Noticeably, all the frontiers by RE in Figure 1b are even inferior to that by GO. As opposed to the wide concern about the ineffectiveness of the well-known resampled mean-variance portfolio in improving Markowitz's portfolio (Wolf 2013), RE in most cases outperforms GO. Third, Figures 1c, 1g and 1k show that for a fixed number of resampled subsets n_3 in the inner loop, there exists a fine structure in the size of each subset n_4. The location or equivalently the performance of the actual frontiers by SS shows a typical U-shape manner in classical Vapnik-Chervonenkis (VC) analysis. This observation implies some tradeoff between diversification benefits and estimation errors, i.e., that the cases with $n_4 = n^{0.5}$ and $n_4 = n^{0.9}$ both underperform the case with $n_4 = n^{0.7}$. Fourth, Figures 1d, 1h and 1l saliently show that SS is more sensitive to the size of each subset n_4 than the number of resampled subsets n_3. In each of those figures the frontiers by SS with the same n_4 are identical, whereas across those figures the corresponding frontiers by SS with the different n_4 have gaps.

From Figure 1m to Figure 1p, we explore the characteristics of EGO. First, the effectiveness of the inner loop chiefly determines the level of performance of EGO, i.e., that all the frontiers by EGO resemble those by SS and dominate those by RE. Second, in line with our observations in the previous figures, although all the four hyperparameters impact EGO, the size of each subset n_4 is most influential. Not only wider gaps among frontiers from different n_4 can be observed, but also the magnitude of the associated gaps is larger than that due to the changes of other hyperparameters. Thus, the study suggests that users of EGO should focus on testing its performance by tuning the subset size n_4 and possibly choose large values for n_1, n_2 and n_3 to increase diversification within a given bound of computing power. Besides cross-validation, we offer one heuristic guideline of selecting n_4: Users could start with the value of n_4 that gives no fewer than five data points per asset and covers 30% to 50% of assets.

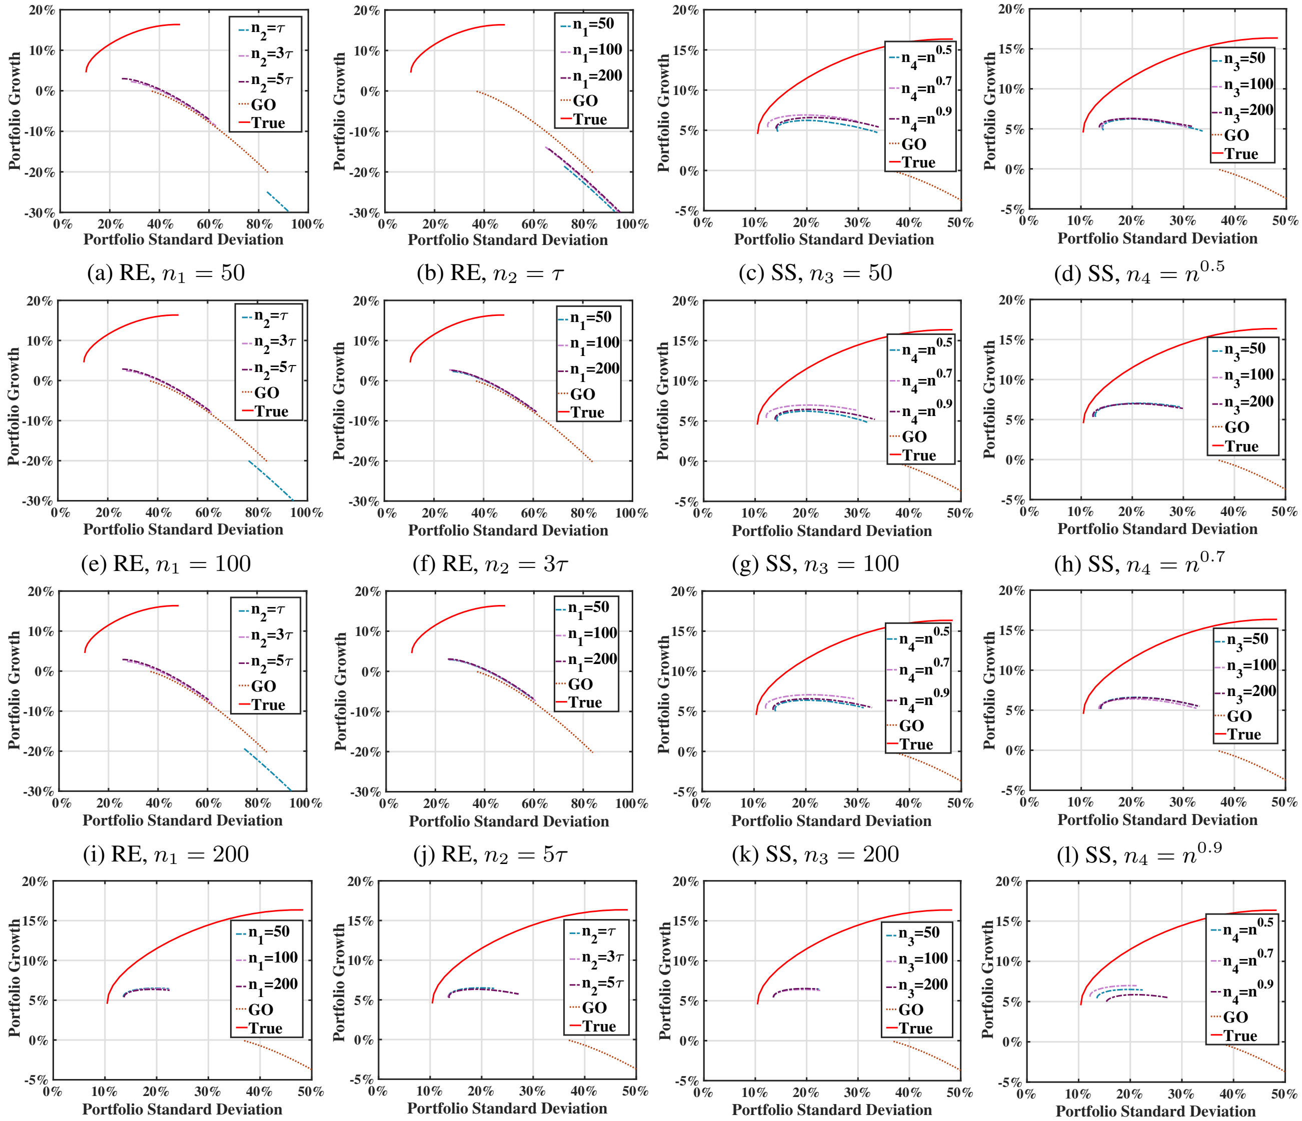

(a) RE, $n_1 = 50$ (b) RE, $n_2 = \tau$ (c) SS, $n_3 = 50$ (d) SS, $n_4 = n^{0.5}$

(e) RE, $n_1 = 100$ (f) RE, $n_2 = 3\tau$ (g) SS, $n_3 = 100$ (h) SS, $n_4 = n^{0.7}$

(i) RE, $n_1 = 200$ (j) RE, $n_2 = 5\tau$ (k) SS, $n_3 = 200$ (l) SS, $n_4 = n^{0.9}$

(m) EGO, Change n_1 (n) EGO, Change n_2 (o) EGO, Change n_3 (p) EGO, Change n_4

Figure 1: Growth frontiers for EGO based on different hyperparameters in a simulation study. The whole dataset FF48 with $n = 48$ assets is used to estimate the means and the covariance matrix as the parameters in a multivariate lognormal return model for simulation. $\tau = 60$ return data for training are then synthesized by simulation. In the last row for EGO, the base hyperparameters are $n_1 = 50$, $n_2 = \tau$, $n_3 = 50$ and $n_4 = n^{0.5}$. One of the four hyperparameters varies in each case.

4 Experiments

In this section, we first elaborate on tested datasets, baseline portfolios and evaluation metrics. Then we conduct comparison studies and report experimental results.

Data

We select four diverse datasets from two categories to fairly appraise the new strategy. The Fama and French datasets as the first category have been recognized as high-quality and standard protocols in portfolio research (Fama and French 1992). Based on different financial segments of the U.S. stock market, the datasets encompass a wide range of indices constructed by historical data. They extensively cover various assets classes and span long periods. FF25 includes monthly returns of 25 portfolios formed on the basis of size and book-to-market ratio over forty years, and FF48 contains monthly returns of 48 portfolios representing different industrial sectors. The second category comprises of the ETF139 and EQ181 datasets downloaded from the stock market. ETF139 has 139 exchange-traded funds as recently popular investments and EQ181 contains 181 equities with the largest market capital from the Russell Top 200 Index.

Table 1 summarizes these two types of testing data. They stress different aspects in performance assessment. On the

Table 1: Summary of the testing datasets

Dataset	Time Period	Frequency	m	n
FF25	07/01/1963 - 12/31/2004	Monthly	498	25
FF48	07/01/1963 - 12/31/2004	Monthly	498	48
ETF139	01/01/2008 - 10/30/2012	Weekly	252	139
EQ181	01/01/2008 - 10/30/2012	Weekly	252	181

one hand, as representative academic benchmarks, the Fama and French datasets highlight the long-term performance with less selection bias. They also cover multiple recessions in their time spans, such as those in 1987, 1997 and 2000. On the other hand, the two real-world datasets reflect the fluctuating market condition after the financial crisis of 2008. By empirical studies on those datasets, we can more confidently understand the performance of each strategy.

Baseline Portfolio Strategies

To scrutinize the proposed strategy EGO and its two degenerated variants RE and SS, we consider another eight competing portfolios in our comparison studies: (a) *Growth optimal portfolio* (GO) as the classical Kelly growth optimal portfolio should inevitably serve as the first benchmark. (b) *Equally-weighted portfolio* (EW) is a naive yet robust strategy. It has shown superior performance among 15 sophisticated models across eight datasets (DeMiguel, Garlappi, and Uppal 2009). (c) *Value-weighted portfolio* (VW) is a passive market mimicking portfolio. Market mimicking portfolios generally surpass a majority of active mutual fund managers in finance industry (Fama and French 2010). (d) *Two-fund portfolio* by (Tu and Zhou 2011) (TZT) as a portfolio blending model mixes the classical mean-variance and EW portfolio for achieving both estimation error cancellation and wealth growth. (e) *Three-fund portfolio* by (Kan and Zhou 2007) (KZT) aggregates the risk-free asset, the classical mean-variance and minimum-variance portfolio to cure the instability in the standard mean-variance framework. (f) *Four-fund portfolio* by (Tu and Zhou 2011) (TZF) further blends KZT and EW to achieve better performance. (g) *Two-fund portfolio* by (Kan, Wang, and Zhou 2016) (KWZ) as an updated version of TZT forms portfolios solely with risky assets. It targets at outperforming EW. (h) *Online passive aggressive mean reversion portfolio* by (Li et al. 2012) (PAMR) is a portfolio strategy via online learning. It has been shown robustly beating 12 portfolio strategies on six datasets. These baselines are state-of-the-art strategies in both finance and machine learning on the relevant topic. For example, TZT, KZT, TZF and KWZ embody up-to-date efforts in multiple fund portfolio aggregating strategies. They share the fundamental thoughts with EGO about estimation risk mitigation via basis portfolio aggregation.

Performance Metrics

In the experiments we employ four standard metrics in finance to justify the performance of each portfolio (Brandt 2010). First, *Sharpe ratio* (SR) as the most broadly adopted risk-adjusted return measure for a portfolio strategy is provided. Without a risk-free asset, SR is calculated as the portfolio return normalized by its standard deviation: SR $= \tilde{r}/\tilde{\sigma}$ with $\tilde{r}$ as the mean of portfolio net returns and $\tilde{\sigma}$ as the corresponding standard deviation:

$$\tilde{r}=\frac{1}{m-\tau}\sum_{k=\tau+1}^{m}\hat{r}_k \text{ and } \tilde{\sigma}=\sqrt{\frac{1}{m-\tau}\sum_{k=\tau+1}^{m}(\tilde{r}-\hat{r}_k)^2}. \quad (3)$$

To compare portfolios with different rebalancing frequencies, we report the annualized Sharpe ratio as $\sqrt{H}$SR, where the scaling factor H is the number of rebalancing times per year. We set $H = 12$ and 52 for monthly and weekly rebalances, respectively. Second, *Volatility* (VO) as a basic quantitative risk measure in finance is offered to gauge risk for each portfolio, i.e., $\tilde{\sigma}$. Likewise, we report the annualized volatility as $\sqrt{H}\tilde{\sigma}$. Third, *Turnover rate* (TO) measuring the volume of transaction and thereby reflecting the impact from market frictions is computed as

$$\text{TO} = \frac{1}{m-\tau}\sum_{k=\tau}^{m-1}||\hat{\boldsymbol{\omega}}_{k^+}-\hat{\boldsymbol{\omega}}_k||_1, \quad (4)$$

where $\hat{\boldsymbol{\omega}}_{k^+}$ denotes the portfolio weight vector before rebalancing at t_{k+1} and $\|\cdot\|_1$ is l_1-norm. The above equation calculates an average absolute value of the rebalancing trades across all the assets and over all the trading periods. A high TO could lead to extra trading costs and drastically degrade after-cost return performance. Fourth, *Maximum drawdown* (MDD) as the maximum percentage drop of the cumulative wealth over a tested time period is computed as

$$\text{MDD} = \max_{k\in[\tau,m]}(1-W_k/M_k), \quad (5)$$

with the running maximum of the cumulative wealth M_k and the cumulative wealth W_j at time t_j obtained by

$$M_k = \max_{j\in[\tau,k]} W_j \text{ with } W_j = W_\tau\prod_{h=\tau+1}^{j}(1+\hat{r}_h). \quad (6)$$

MDD is one topmost risk measure for fund managers, because large drawdowns often trigger fund redemptions.

In addition, for each strategy in our calculation we apply the "rolling-horizon" setting for the sequential out-of-sample performance evaluation (DeMiguel, Garlappi, and Uppal 2009). In particular, from time t_τ to t_m, at each rebalancing time t_k for $k = \tau, \ldots, m$, we first calculate the portfolio weight $\hat{\boldsymbol{\omega}}_k$ based on the return data $\{\mathbf{R}_l\}_{l=k-\tau+1}^{k}$. Then, we compute the realized out-of-sample net return $\hat{r}_k$ for the subsequent trading period. Then, we evaluate the out-of-sample characteristics of portfolios in the discussed four standard metrics by the achieved sequences of $\hat{r}_k$ and $\hat{\boldsymbol{\omega}}_k$.

Results

Across the four datasets Table 2 reports the overall performance of the compared 11 portfolios, including the proposed EGO and its two degenerated variants SS and RE. First, EGO outruns other strategies in most cases. A progressive performance enhancement is demonstrated from GO to EGO. In particular, RE is slightly superior to GO, SS mostly

Table 2: Portfolio performance of strategies

Dataset	Metrics	EGO	SS	RE	GO	EW	VW	TZT	KZT	TZF	KWZ	PAMR
FF25	SR	**1.11**	1.09	0.71	0.70	0.80	0.87	-0.03	0.20	0.34	0.53	0.94
	p-value	0.01	0.02	0.66	0.65	1.00	0.01	0.31	0.02	0.15	0.27	0.27
	VO (%)	**13.79**	13.93	27.51	27.68	17.84	17.67	109.98	881.12	53.14	42.44	19.91
	TO (%)	85.10	138.11	1087.23	1100.42	1.81	**0.00**	2572.46	23936.66	617.56	1930.11	0.01
	MDD (%)	**36.82**	37.35	47.51	47.44	47.77	47.81	-100.00	-100.00	-100.00	86.80	54.46
FF48	SR	**1.02**	1.00	0.44	0.43	0.71	0.73	0.25	0.27	0.75	0.46	0.69
	p-value	0.05	0.05	0.24	0.22	1.00	0.60	0.03	0.03	0.49	0.29	0.28
	VO (%)	**12.37**	13.20	25.32	25.25	17.07	16.42	85.39	1000.95	17.46	33.22	18.57
	TO (%)	33.01	82.13	463.92	462.63	3.14	**0.00**	1435.98	37083.69	100.82	759.08	0.24
	MDD (%)	**34.91**	35.80	75.92	76.15	48.33	45.74	-100.00	-100.00	43.08	66.7	47.43
ETF139	SR	**1.45**	1.40	1.45	1.33	0.51	0.50	0.42	0.97	0.32	-0.27	0.86
	p-value	0.31	0.34	0.26	0.35	1.00	0.55	0.82	0.14	0.51	0.19	0.21
	VO (%)	**4.11**	4.40	5.55	5.62	21.20	21.06	21.78	56.46	18.6	22.57	36.85
	TO (%)	89.35	273.15	286.86	357.52	1.08	**0.00**	403.91	10520.42	338.78	5762.82	153.38
	MDD (%)	**2.02**	2.38	5.35	5.56	22.35	22.51	24.83	19.35	20.44	28.11	34.33
EQ181	SR	1.41	1.34	0.52	0.67	0.97	0.97	1.34	-0.24	1.00	**1.96**	0.30
	p-value	0.77	0.82	0.75	0.83	1.00	0.84	0.41	0.41	0.87	0.43	0.21
	VO (%)	**8.58**	8.87	24.16	23.94	15.43	15.29	17.45	45.50	15.07	45.57	22.86
	TO (%)	33.90	99.17	580.96	587.33	1.85	**0.00**	140.72	1332.92	55.33	114.74	140.48
	MDD (%)	**4.84**	5.67	9.54	9.53	9.19	9.17	10.21	22.18	8.45	16.81	18.70

Note: We use short sliding widows $\tau = 30, 60, 150$ and 200 for the four datasets, respectively. EGO, SS, RE and GO implement the growth optimal portfolio with the return target as the average of minimum and maximum estimated returns at each step. The relevant parameters for EGO, SS and RE are from $n_1 = 100$, $n_2 = 5\tau$, $n_3 = 100$ and $n_4 = n^{0.7}$. The risk-aversion coefficient in TZT, KZT, TZF and KWZ is three as the original papers. The p-values under SR quantify the statistical significance of the difference in SR between two portfolios with EW as the benchmark portfolio. As portfolio returns are seldom i.i.d., the studentized circular block bootstrapping method is used to compute the p-values with 5000 bootstrap resamples and a block with the size of five (Ledoit and Wolf 2008).

outperforms both RE and GO at a marked effect size, and EGO further improves SS with the aid of its outer loop for bagging. Those observations that RE exceeds GO and EGO surpasses SS also imply that bagging, whose usefulness to Markowitz's portfolio becomes debatable (Wolf 2013), is effectual in improving Kelly's portfolio. That may indicate that bagging helps by keeping a lid on the exacerbating moments estimation errors in Kelly's portfolio.

Second, VW and EW as two passive baselines produce highly competitive results. Consonant with the observed robustness in (DeMiguel, Garlappi, and Uppal 2009), VW and EW in this study also largely beat all strategies other than EGO and SS. In additional, their TOs are artificially low due to their intrinsic design for marginal transaction.

Third, TZT, KZT, TZF and KWZ as the exclusively designed multiple fund aggregating strategies underperform EGO, SS and others in most cases. Although these four strategies share the similar thoughts of canceling estimation errors by aggregating multiple portfolios as EGO and SS, their original goal can hardly be realized in the condition with a small training sample. It is unsurprising because many of their desired properties are only asymptotically achievable when the sample size is large. Because of the paucity of training data, the estimation risk in those strategies overwhelms the benefits from peculiarly designed portfolio structure. As such, they all generate high TOs and in a few cases are even ruined, i.e., MDD $= -100\%$. Portfolios with large estimation errors in input parameters often show high turnovers due to huge variation of the optimal solution. Theoretically, when the training sample is small, out-of-sample performance will be substantially penalized by model complexity, i.e., that the sample complexity for a complex model is high. That has occurred to the preceding four strategies. In contrast, EGO and SS without having stringent requirements on asymptotic conditions diminish the model complexity for a fixed sample size, thereby lowering the sample complexity and enhancing out-of-sample performance. Thus, EGO and SS reap benefits from error cancellation more than the four aggregating strategies.

Fourth, PAMR is superior to TZT, KZT, TZF and KWZ, and is comparable to EW and VW but underperforms EGO and SS. In view of the above, when data are insufficient to build complex but robust strategies, besides EGO and SS investors may also consider online strategies.

5 Conclusions and Discussions

In this paper, we have presented a new portfolio construction strategy via ensemble learning. It has filled some voids of empirical studies on the classical Kelly growth optimal portfolio. Underlining the challenging but realistic condition with noisy small samples, we have analyzed the hyperparameter selection and compared the proposed portfolio with 10 other strategies on four diversified datasets. In extensive evaluation criteria, our studies have demonstrated that the proposed portfolio outperforms peers in most cases. We believe that it represents a fresh effort in hastening the applicability of machine learning algorithms to finance research and will cross-fertilize ideas and techniques in specific topics, such as enhancing stock prediction accuracy via incorporating texts from social media (Wu et al. 2018).

References

Agarwal, A.; Hazan, E.; Kale, S.; and Schapire, R. E. 2006. Algorithms for portfolio management based on the Newton method.

In *Proceedings of the 23rd International Conference on Machine Learning*, 9–16.

Bottazzi, G., and Santi, C. 2017. Relative performance of mean-variance, Kelly and universal portfolios in the equity market. Working Paper.

Brandt, M. W. 2010. Portfolio choice problems. In Ait-Sahalia, Y., and Hansen, L. P., eds., *Handbooks of Financial Econometrics*. Elsevier. 269–336.

Breiman, L. 1961. Optimal gambling systems for favorable games. In *Proceedings of the 4th Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, 65–78.

Breiman, L. 1996. Bagging predictors. *Machine Learning* 24(2):123–140.

Broadie, M. 1993. Computing efficient frontiers using estimated parameters. *Annals of Operations Research* 45(1):21–58.

Christensen, M. 2005. On the history of the growth optimal portfolio. Working Paper.

Cover, T. M. 1991. Universal portfolios. *Mathematical Finance* 1(1):1–29.

DeMiguel, V.; Garlappi, L.; and Uppal, R. 2009. Optimal versus naive diversification: How inefficient is the $1/N$ portfolio strategy? *Review of Financial Studies* 22:1915–1953.

Derbeko, P.; El-Yaniv, R.; and Meir, R. 2002. Variance optimized bagging. In *Proceedings of the 13th European Conference on Machine Learning*, 60–71.

Estrada, J. 2010. Geometric mean maximization: An overlooked portfolio approach? *Journal of Investing* 19(4):134–147.

Fama, E. F., and French, K. R. 1992. The cross-section of expected stock returns. *Journal of Finance* 47(2):427–465.

Fama, E. F., and French, K. R. 2010. Luck versus skill in the cross-section of mutual fund returns. *Journal of Finance* 65(5):1915–1947.

Györfi, L.; Ottucsák, G.; and Walk, H. 2012. *Machine Learning for Financial Engineering*, volume 8. World Scientific.

Ho, T. K. 1998. The random subspace method for constructing decision forests. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20(8):832–844.

Hunt, B. F. 2005. Feasible high growth investment strategy: Growth optimal portfolios applied to Dow Jones stocks. *Journal of Asset Management* 6(2):141–157.

Jagannathan, R., and Ma, T. 2003. Risk reduction in large portfolios: Why imposing the wrong constraints helps. *Journal of Finance* 58(4):1651–1683.

Kalai, A., and Vempala, S. 2002. Efficient algorithms for universal portfolios. *Journal of Machine Learning Research* 3:423–440.

Kan, R., and Zhou, G. 2007. Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis* 42(03):621–656.

Kan, R.; Wang, X.; and Zhou, G. 2016. Optimal portfolio selection with and without risk-free asset. *Available at SSRN 2729429*.

Kelly, J. L. 1956. A new interpretation of information rate. *Bell Labs Technical Journal* 35(4):917–926.

Kleiner, A.; Talwalkar, A.; Sarkar, P.; and Jordan, M. I. 2014. A scalable bootstrap for massive data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 76(4):795–816.

Kolm, P. N.; Tütüncü, R.; and Fabozzi, F. J. 2014. 60 years of portfolio optimization: Practical challenges and current trends. *European Journal of Operational Research* 234(2):356–371.

Latané, H. A. 1959. Criteria for choice among risky ventures. *Journal of Political Economy* 67(2):144–155.

Ledoit, O., and Wolf, M. 2008. Robust performance hypothesis testing with the Sharpe ratio. *Journal of Empirical Finance* 15:850–859.

Li, B., and Hoi, S. C. 2015. *Online Portfolio Selection: Principles and Algorithms*. CRC Press.

Li, B.; Hoi, P.; Hoi, S. C.; and Gopalkrishnan, V. 2012. PAMR: Passive aggressive mean reversion strategy for portfolio selection. *Machine Learning* 87:221–258.

Li, B.; Wang, J.; Huang, D.; and Hoi, S. C. 2017. Transaction cost optimization for online portfolio selection. *Quantitative Finance* 1–14.

Luenberger, D. G. 1998. *Investment Science*. Oxford University Press.

MacLean, L. C.; Thorp, E. O.; and Ziemba, W. T. 2011. *The Kelly Capital Growth Investment Criterion: Theory and Practice*, volume 3. World Scientific.

Markowitz, H. 1952. Portfolio selection. *Journal of Finance* 7(1):77–91.

Michaud, R. O. 1989. The Markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts Journal* 45(1):31–42.

Polikar, R. 2006. Ensemble based systems in decision making. *IEEE Circuits and Systems Magazine* 6(3):21–45.

Rokach, L. 2010. Ensemble-based classifiers. *Artificial Intelligence Review* 33(1-2):1–39.

Samuelson, P. A. 1979. Why we should not make mean log of wealth big though years to act are long. *Journal of Banking & Finance* 3(4):305–307.

Shen, W., and Wang, J. 2016. Portfolio blending via Thompson sampling. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, 1983–1989.

Shen, W., and Wang, J. 2017. Portfolio selection via subset resampling. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 1517–1523.

Shen, W.; Wang, J.; Jiang, Y.-G.; and Zha, H. 2015. Portfolio choices with orthogonal bandit learning. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, 974–980.

Skurichina, M., and Duin, R. P. W. 2002. Bagging, boosting and the random subspace method for linear classifiers. *Pattern Analysis & Applications* 5(2):121–135.

Tu, J., and Zhou, G. 2011. Markowitz meets Talmud: A combination of sophisticated and naive diversification strategies. *Journal of Financial Economics* 99:204–215.

Uziel, G., and El-Yaniv, R. 2018. Growth-optimal portfolio selection under CVaR constraints. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*, 48–57.

Wolf, M. 2013. Resampling vs. shrinkage for benchmarked managers. *Wilmott Magazine* 76–81.

Wu, H.; Zhang, W.; Shen, W.; and Wang, J. 2018. Hybrid deep sequential modeling for social text-driven stock prediction. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, 1627–1630.

Zhou, Z.-H. 2012. *Ensemble Methods: Foundations and Algorithms*. CRC press.

Ziemba, W. T. 2015. A response to Professor Paul A. Samuelson's objections to Kelly capital growth investing. *Journal of Portfolio Management* 42(1):153–167.