Resolution of degeneracy in Merton's portfolio problem

Pun, Chi Seng; Wong, Hoi Ying

2016

Pun, C. S., & Wong, H. Y. (2016). Resolution of degeneracy in merton's portfolio problem. SIAM Journal on Financial Mathematics, 7(1), 786-811. doi:10.1137/16M1065021

<https://hdl.handle.net/10356/89935>

<https://doi.org/10.1137/16M1065021>

© 2016 Society for Industrial and Applied Mathematics. This paper was published in SIAM Journal on Financial Mathematics and is made available as an electronic reprint (preprint) with permission of Society for Industrial and Applied Mathematics. The published version is available at: [<http://dx.doi.org/10.1137/16M1065021>]. One print or electronic copy may be made for personal use only. Systematic or multiple reproduction, distribution to multiple locations via electronic or other means, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper is prohibited and is subject to penalties under law.

Downloaded on 26 Apr 2023 20:12:06 SGT

Resolution of Degeneracy in Merton's Portfolio Problem*

Chi Seng Pun[†] and Hoi Ying Wong[‡]

Abstract. The Merton problem determines the optimal intertemporal portfolio choice by maximizing the expected utility and is the basis of modern portfolio theory in continuous-time finance. However, its empirical performance is disappointing. The estimation errors of the expected rates of returns make the optimal policy degenerate, resulting in an extremely low (or unbounded) expected utility value for a high-dimensional portfolio. We further prove that the estimation error of the variance-covariance matrix leads to the degenerated policy of solely investing in the risk-free asset. This study proposes a constrained ℓ_1 -minimization approach to resolve the degeneracy in the high-dimensional setting and stabilize the performance in the low-dimensional setting. The proposed scheme can be implemented with simple linear programming and involves negligible additional computational time, compared to standard estimation. We prove the consistency of our framework that our estimate of the optimal control tends to be the true one. We also derive the rate of convergence. Simulation studies are provided to verify the finite-sample properties. An empirical study using S&P 500 component stock data demonstrates the superiority of the proposed approach.

Key words. high-dimensional portfolio, Merton's problem, expected utility maximization, constrained ℓ_1 -minimization, Dantzig selector, sparsity

AMS subject classifications. 91G10, 91G70, 93E20

DOI. 10.1137/16M1065021

1. Introduction. The seminal papers by Merton (1969 and 1971), reprinted in [33], pioneered the intertemporal portfolio choice using the expected utility maximization (EUM) in continuous time. The Merton portfolio problem has been the benchmark for investigating economic behavior and investment demands under different market conditions. However, when optimized portfolios, such as the Merton strategy and others, are tested with real data, their empirical performance is far from satisfactory.

The crux of the problem is the estimation errors of the model parameters. It is shown in [18, 24] that plugging the sample estimates of the expected rates of returns (drifts) into the optimal strategies obtained by maximizing the expected logarithmic and power utilities leads both strategies to degenerate to solutions leading to an extremely low expected utility value. The expected utility may tend toward negative infinity in a precise technical term as the number of risky assets increases, even though the variance-covariance matrix of the asset

*Received by the editors March 9, 2016; accepted for publication (in revised form) September 13, 2016; published electronically November 15, 2016.

<http://www.siam.org/journals/sifin/7/M106502.html>

Funding: The second author was supported by the Research Grants Council of Hong Kong, project GRF2130469.

[†]Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore (cspon@ntu.edu.sg).

[‡]Corresponding author. Department of Statistics, Chinese University of Hong Kong, Shatin, Hong Kong, China (hywong@cuhk.edu.hk).

returns is known in advance [18]. In a single-period quadratic program with linear constraints, sample estimates cause the variance of the optimal policy to be underestimated [19].

One immediate method of reducing the estimation errors is to use a long time series or high-frequency data. However, estimating the drift alone requires more than a hundred years of data to achieve a satisfactory result [32], and the accuracy of the mean estimate is independent of the sampling frequency [18]. The estimation of the covariance matrix for high-frequency data requires additional treatment [2]. A long estimation window also increases the likelihood of nonstationarity in the estimated parameters [7].

Another remedy for reducing the estimation errors of the mean vector is to use biased or shrinkage estimators, pioneered by Stein, James, and Bock [5, 28, 36]. In the context of portfolio analysis, Jorion [29] suggested using the James–Stein estimator to shrink the sample mean toward the mean of the minimum-variance portfolio. However, the comprehensive empirical study in [16] shows that the James–Stein estimator does not improve the portfolio performance significantly.

Recent improvements in theoretical optimal portfolios impose norm constraints on the portfolio weights of the original optimization problem. The intention is to limit the gross exposure in risky assets and therefore reduce the effect of the accumulated estimation errors. Similar ideas are found in [8, 15, 27]. When the ℓ_1 -norm constraints are applied to the portfolio weights, as in [18, 24], they can be efficiently implemented in the fashion of least absolute shrinkage and selection operator (LASSO). In particular, the ℓ_1 -norm constraint eventually binds the objective value so it does not degenerate to negative infinity [18]. The norm-constrained approach is, however, restricted to static optimization problems and/or constant rebalancing portfolios. From an optimization perspective the original Merton portfolio, which contains no norm constraints, has a larger opportunity set and the solution should perform better than that with additional constraints, once the model parameter values are known. The variance of the ℓ_1 -norm-constrained optimal portfolio converges to that of the Merton strategy under the single period setting [20], but it is still unclear if the constrained strategy asymptotically converges fully on the underlying true Merton strategy. Our newly proposed implementation ensures that the estimated policy asymptotically converges to the true Merton strategy which is free of estimation errors.

The drifts and the variance-covariance matrix (hereafter referred to as the covariance matrix) among risky asset returns both suffer from serious estimation errors. The latter hugely distorts the discrete-time mean-variance strategies to become a random stock picking scheme [35]. This study shows the effect of the accumulated estimation errors in the covariance matrix for Merton's problem. To stabilize the variance estimate, a practical solution is to use information from implied volatilities. Therefore, the EUM problem with a stochastic volatility has been developed in [3, 22, 38], but implied volatility data is available only for some risky assets and major financial indices and the estimation of correlations is still a major concern. An optimal portfolio robust to correlation was recently established in [21], but its extension to a high-dimensional portfolio is nontrivial. The implementation proposed in this study is particularly useful in resolving the degeneracy associated with estimation errors in the covariance matrix.

In this study, we propose a constrained ℓ_1 -minimization approach to directly estimate the effective parameters in the EUM portfolio for a general class of utility functions, including

exponential, power, and mixture of powers utilities. The constrained ℓ_1 -minimization is widely used in high-dimensional statistical problems such as [10, 11, 13]. It not only can address the curse of dimensionality but also performs well in a low-dimensional setting. Its development is strongly influenced by the theoretical work of Candès and Tao [13] and the series of discussion papers attached to this publication, such as [12].

We comprehensively investigate the constrained ℓ_1 -minimization for the Merton portfolio problem. We derive the oracle property of our proposed estimator for $(\log p)/n \rightarrow 0$, where p is the number of risky assets and n is the number of observation time points. The oracle property means that our approach asymptotically recovers the ideal optimal strategy which is free of estimation error. The oracle property holds even when both estimation errors in the drifts and the covariance matrix exist. To the best of our knowledge, this newly proposed implementation is the first unified framework for the Merton problem that works for $p/n \rightarrow 0$, a constant or infinity (with a reasonable growth in p relative to n). Although we face situations with a large p in practice, the ratio between p and n can have many possibilities. Yet, our approach is sufficient for most practical situations about the ratio of p and n . We compare portfolio performance with plug-in portfolios and norm-constrained portfolios numerically and empirically.

We would like to highlight some important contributions of this paper in mathematics and finance. It has been shown in [18, 24] that the total expected utility value degenerates to negative infinity due to the estimation errors in the drifts for a known covariance matrix and $p/n \rightarrow \infty$. We further clarify the degeneracy of the EUM problem caused by estimation errors in the drifts and the covariance matrix. We prove that the expected utility conditional on the historical data degenerates to negative infinity in probability for several widely used utility functions once the drifts are unknown parameters for $p/n \rightarrow \infty$. Our result is stronger because the degeneracy in conditional expectation implies the degeneracy in the total expectation, but the reverse is generally not true. We also prove the degeneracy of the conditional expected utility associated with the estimation errors in the covariance matrix for $p/n \rightarrow \infty$.

While the first contribution is more toward mathematics, the second one is relevant to finance. We propose a new implementation of the EUM portfolios, which is realized with linear programming, and hence call the resulting portfolio the linear programming optimal (LPO) portfolio. The LPO portfolio works well for both low- and high-dimensional settings. More precisely, we mathematically prove the oracle property, in Theorem 4.2, that the LPO approach asymptotically recovers the true EUM portfolio once $(\log p)/n \rightarrow 0$. In other words, our approach is generally applicable for p/n tending to zero, a constant, and infinity. Our numerical studies verify the theoretical results and demonstrate the finite-sample properties with the mixture of power utilities, and the norm-constrained strategy for this is not currently available in the literature.

We also offer empirical evidence to the application of the LPO approach. Using the dataset of S&P500 component stocks, we conduct an out-of-sample empirical study to compare the equally weighted (EW) portfolio, the plug-in EUM portfolio using sample/James–Stein estimates, the norm-constrained EUM portfolio, and our LPO portfolio. We also compare the ability of the norm-constrained and the LPO approaches to screen assets. Most results favor the LPO portfolios or the LPO-based portfolios, especially when the portfolio turnovers are taken into account.

The rest of this paper is organized as follows. Section 2 reviews the EUM, sets the notation, and describes the optimal portfolio. Section 3 mathematically proves the pitfalls of the plug-in EUM portfolios for two cases: the covariance matrix is known but the drifts are unknowns, and the drifts are known but the covariance matrix is an unknown. Section 4 presents the new approach in implementing the EUM portfolios and proves its oracle properties. Section 5 verifies our assertions and examines finite-sample properties under different correlation structures among risky assets through simulation studies. Empirical studies in section 6 compare the proposed LPO approach against other existing alternatives using the dataset of S&P500 component stocks. Section 7 concludes the paper.

2. The Merton portfolio problem. Consider the complete filtered (physical) probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \in [0, T]}, \mathbb{P})$, where T is a constant investment horizon and the filtration $\{\mathcal{F}_t\}_{t \in [0, T]}$ is generated by the p -dimensional standard Brownian motion $W(t) = (W_1(t), \dots, W_p(t))$. We denote by $\mathcal{L}_{\mathcal{F}}^2(0, T; \mathbb{R}^p)$ the set of all $\{\mathcal{F}_t\}_{t \geq 0}$ -adapted $\mathbb{R}^p$ -valued measurable stochastic processes $a(t)$ such that $\mathbb{E}[\int_0^T \|a(t)\|_2^2 dt] < \infty$.

2.1. Financial market. Consider a securities market with p risky assets (stocks) and one risk-free asset (bond) traded continuously in time. The price processes of the risky assets $\{S_i(t)\}_{i=1}^p$ evolve as follows:

$$(2.1) \quad \begin{cases} dS_i(t) &= S_i(t)[\mu_i dt + \sigma_i dW_t], \quad t \in [0, T], \\ S_i(0) &= s_i > 0, \end{cases}$$

where μ_i is the appreciation rate and $\sigma_i = (\sigma_{i1}, \dots, \sigma_{ip}) \in \mathbb{R}^p$ represents the volatility of the i th stock, for $i = 1, \dots, p$. The risk-free asset price $S_0(t)$ satisfies

$$\begin{cases} dS_0(t) &= r_f S_0(t) dt, \quad t \in [0, T], \\ S_0(0) &= s_0 > 0, \end{cases}$$

where $r_f > 0$ is the constant risk-free interest rate.

We consider an investor with a self-financing portfolio $\{N_i(t)\}_{i=0}^p$ whose total wealth at time t is denoted by $X(t)$, where $N_i(t)$ is the number of shares of the i th asset invested at time t . Then

$$X(t) = \sum_{i=0}^p N_i(t) S_i(t) = \sum_{i=0}^p u_i(t), \quad t \in [0, T],$$

where $u_i(t) := N_i(t) S_i(t)$ is the amount of money invested in the i th asset at time t and the remaining wealth held in the risk-free asset. Ignoring the transaction costs and consumptions, the dynamics of the wealth process $X(t)$ is given by

$$(2.2) \quad \begin{cases} dX(t) &= [r_f X(t) + u(t)' \beta] dt + u(t)' \sigma dW_t, \quad t \in [0, T], \\ X(0) &= x_0 > 0, \end{cases}$$

where $u(t) = (u_1(t), \dots, u_p(t))'$ describes the portfolio of the investor, $\beta = (\mu_1 - r_f, \dots, \mu_p - r_f)'$ is the excess return vector, and $\sigma = (\sigma'_1, \dots, \sigma'_p)' \in \mathbb{R}^{p \times m}$ is the volatility matrix of the p risky assets. We also assume the following nondegeneracy condition for the covariance matrix (Σ) of the risky assets is satisfied:

$$(2.3) \quad \Sigma := \sigma \sigma' \succeq \delta I \text{ for some } \delta > 0,$$

where $A \succeq B$ for matrices A and B means that $A - B$ is nonnegative-definite.

2.2. The optimal portfolios. To investigate the degeneracies of the Merton problem with general utility functions, we review the optimal policy in these terms and standardize the notation. Consider a utility function $U : \mathbb{R} \rightarrow \mathbb{R}$ which is C^1 , increasing, and strictly concave and satisfies Inada conditions (see [26]):

$$U'(-\infty) = +\infty, \quad U'(+\infty) = 0.$$

If $U : \mathbb{R}_+ \rightarrow \mathbb{R}$, such as power utility and mixture of power utilities, the Inada condition $U'(-\infty) = +\infty$ is replaced by $U'(0+) = +\infty$ and the asymptotic elasticity condition is imposed: $AE[U] := \lim_{x \rightarrow \infty} x \frac{U'(x)}{U(x)} < 1$. These are the usual conditions for utility functions. For these two types of utility functions, the following analyses hold.

Typical EUM problems define the value function of the investor as

$$V(t, x) = \sup_{u(t) \in \mathcal{L}_{\mathcal{F}}^2(0, T; \mathbb{R}^p)} \mathbb{E}^{\mathbb{P}}[U(X^u(T)) | X(t) = x],$$

where $X^u(T)$, driven by (2.2), is the wealth of the investor who adopts the portfolio u . If $U : \mathbb{R}_+ \rightarrow \mathbb{R}$, an admissible control $u(t)$ is required to almost completely ensure the positivity of the wealth, i.e., $X^u(t) > 0$, a.s., $t \in [0, T]$. This should be kept in mind when we investigate the execution of a portfolio. The EUM problems are well-developed in the literature. The optimal EUM portfolio takes the form

$$(2.4) \quad u^*(t, x) = R(t, x) \Sigma^{-1} \beta,$$

where $R(t, x) := -V_x(t, x)/V_{xx}(t, x)$ is the risk-tolerance function. As derived in [22] and [34], this function solves the following nonlinear partial differential equation (PDE):

$$R_t + \frac{\Theta_p}{2} R^2 R_{xx} + r_f x R_x - r_f R = 0, \quad R(T, x) = -U'(x)/U''(x),$$

where $\Theta_p = \beta' \Sigma^{-1} \beta$. For a general utility function, the above PDE can be solved only numerically. We can also numerically solve for V with the large wealth asymptotic results:

$$V_x(t, x) = V_x(t, x^*) \exp \left(\int_x^{x^*} \frac{1}{R(t, y)} dy \right), \quad \text{and} \quad V(t, x) = V(t, x^*) - \int_x^{x^*} V_x(t, y) dy,$$

where x^* is large enough that we can use large wealth asymptotics to insert $V_x(t, x^*)$ and $V(t, x^*)$.

The risk tolerance function and the value function admit explicit expressions for specific utility functions. The analytical tractability enables us to draw economic interpretations. For the log utility $U_{\log}(x) = \log x$, we have $R_{\log}(t, x) = x$ and

$$(2.5) \quad u_{\log}^*(t, x) = x \Sigma^{-1} \beta,$$

$$(2.6) \quad V_{\log}(t, x) = \left(r_f + \frac{\Theta_p}{2} \right) (T - t);$$

for the power (constant relative risk aversion (CRRA)) utility $U_{pow}(x) = \frac{x^{1-\gamma}}{1-\gamma}$, $\gamma \neq 1$, we have $R_{pow}(t, x) = \frac{x}{\gamma}$ and

$$(2.7) \quad u_{pow}^*(t, x) = \frac{x}{\gamma} \Sigma^{-1} \beta,$$

$$(2.8) \quad V_{pow}(t, x) = \frac{x^{1-\gamma}}{1-\gamma} \exp \left\{ \left[r_f(1-\gamma) + \frac{1-\gamma}{2\gamma} \Theta_p \right] (T-t) \right\};$$

for the exponential (constant absolute risk aversion) utility $U_{exp}(x) = -\frac{1}{\gamma} e^{-\gamma x}$, we have $R_{exp}(t, x) = \frac{e^{-r_f(T-t)}}{\gamma}$ and

$$(2.9) \quad u_{exp}^*(t, x) = \frac{e^{-r_f(T-t)}}{\gamma} \Sigma^{-1} \beta,$$

$$(2.10) \quad V_{exp}(t, x) = -\frac{1}{\gamma} \exp \left\{ -\gamma x e^{r_f(T-t)} - \frac{1}{2} \Theta_p (T-t) \right\},$$

where γ is a positive constant that represents the degree of risk aversion.

These explicit formulas assist in this particular study, but our proposed approach is also widely applicable to a general class of utility functions, including the family of mixture of power utilities, for which the corresponding ℓ_1 -norm-constrained optimal policy is not available. We thoroughly examine the implementation of the EUM portfolios associated with the log, power, and exponential utilities in the subsequent section and demonstrate the application of our approach to the mixture of power utilities in numerical studies, under both low- and high-dimensional settings.

3. Degeneracy in Merton's problem. In practice, we need to estimate Σ and β in order to implement the Merton problem. This section investigates the performance of traditional plug-in portfolios, in which β and Σ are replaced with their sample estimates in (2.5), (2.7), and (2.9). The effects of the accumulated estimation errors in the drifts for the log and power utilities are discussed in [24] and [18], respectively, under the assumption that the Σ is known. We establish stronger and novel results for two cases: known Σ and unknown β , and known β and unknown Σ .

In this section, we illustrate the degeneracy in Merton's problem with p -asymptotic results. We show that the out-of-sample value function of an EUM portfolio with plug-in sample mean estimate (for a randomized training dataset) degenerates to the infimum of the corresponding utility function as the number of assets increases. It generalizes the results in [24, 18], which show only the total expectation of utility degenerates. Moreover, we show that even the excess return vector β is observable (free of estimation error), all the plug-in EUM portfolios tend to suggest no investment in risky assets under the estimation error of the covariance matrix Σ . It generally deviates from the investors' expectations.

3.1. Sample estimates of β and Σ . Suppose we have the daily log returns of p risky assets as $\{r^{(l)} = (r_1^{(l)}, \dots, r_p^{(l)})', 1 \leq l \leq n\}$, where n is the number of observations and the time-step is defined as $\delta t = 1/252$ for 252 trading days a year. Denote $D = \sigma(r^{(1)}, \dots, r^{(n)})$

as the σ -field generated by the historical data. Clearly, the maximum likelihood estimators (MLEs) for β and Σ are, respectively,

$$(3.1) \quad \hat{\beta}_n = \hat{\mu}_n - r_f \mathbf{1}_p, \quad \hat{\mu}_n = \frac{1}{\delta t} \bar{r} + \frac{1}{2} \text{diag}(\hat{\Sigma}_n),$$

$$(3.2) \quad \hat{\Sigma}_n = \frac{1}{n\delta t} \sum_{l=1}^n (r^{(l)} - \bar{r})(r^{(l)} - \bar{r})',$$

where $\mathbf{1}_p = (1, \dots, 1)' \in \mathbb{R}^p$ and $\bar{r} = \frac{1}{n} \sum_{l=1}^n r^{(l)}$.

According to the plug-in strategies of (2.5), (2.7), and (2.9) with the MLEs (3.1) and (3.2), the corresponding out-of-sample terminal wealths from (2.2) conditional on information at t are solved by Itô's lemma as

$$\begin{aligned} X^{\widehat{u}_{log}^*}(T) &= x \exp \left\{ \left[r_f + \beta' \hat{\Sigma}_n^{-1} \hat{\beta}_n - \frac{\hat{\beta}_n' \hat{\Sigma}_n^{-1} \Sigma \hat{\Sigma}_n^{-1} \hat{\beta}_n}{2} \right] (T-t) + \hat{\beta}_n' \hat{\Sigma}_n^{-1} \sigma(W_T - W_t) \right\}, \\ X^{\widehat{u}_{pow}^*}(T) &= x \exp \left\{ \left[r_f + \frac{\beta' \hat{\Sigma}_n^{-1} \hat{\beta}_n}{\gamma} - \frac{\hat{\beta}_n' \hat{\Sigma}_n^{-1} \Sigma \hat{\Sigma}_n^{-1} \hat{\beta}_n}{2\gamma^2} \right] (T-t) + \frac{1}{\gamma} \hat{\beta}_n' \hat{\Sigma}_n^{-1} \sigma(W_T - W_t) \right\}, \\ X^{\widehat{u}_{exp}^*}(T) &= x e^{r_f(T-t)} + \frac{T-t}{\gamma} \beta' \hat{\Sigma}_n^{-1} \hat{\beta}_n + \frac{1}{\gamma} \hat{\beta}_n' \hat{\Sigma}_n^{-1} \sigma(W_T - W_t), \end{aligned}$$

and the computation of the corresponding out-of-sample value functions is straightforward:

$$(3.3) \quad \begin{aligned} \widehat{V}_{log}(t, x) &:= \mathbb{E}_D[U_{log}(X^{\widehat{u}_{log}^*}(T))] \\ &= \log x + \left[r_f + \beta' \hat{\Sigma}_n^{-1} \hat{\beta}_n - \frac{\hat{\beta}_n' \hat{\Sigma}_n^{-1} \Sigma \hat{\Sigma}_n^{-1} \hat{\beta}_n}{2} \right] (T-t), \end{aligned}$$

$$(3.4) \quad \begin{aligned} \widehat{V}_{pow}(t, x) &:= \mathbb{E}_D[U_{pow}(X^{\widehat{u}_{pow}^*}(T))] \\ &= \frac{x^{1-\gamma} e^{r_f(1-\gamma)(T-t)}}{1-\gamma} \exp \left\{ \frac{(1-\gamma)(T-t)}{\gamma} \left(\beta' \hat{\Sigma}_n^{-1} \hat{\beta}_n - \frac{\hat{\beta}_n' \hat{\Sigma}_n^{-1} \Sigma \hat{\Sigma}_n^{-1} \hat{\beta}_n}{2} \right) \right\}, \end{aligned}$$

$$(3.5) \quad \begin{aligned} \widehat{V}_{exp}(t, x) &:= \mathbb{E}_D[U_{exp}(X^{\widehat{u}_{exp}^*}(T))] \\ &= -\frac{1}{\gamma} \exp \left\{ -\gamma x e^{r_f(T-t)} - \left(\beta' \hat{\Sigma}_n^{-1} \hat{\beta}_n - \frac{\hat{\beta}_n' \hat{\Sigma}_n^{-1} \Sigma \hat{\Sigma}_n^{-1} \hat{\beta}_n}{2} \right) (T-t) \right\}, \end{aligned}$$

which are the conditional expectations of the utilities given the data set D .

Many empirical studies, such as [14, 17, 32], show that the estimation errors seriously distort the optimality of the Merton problem, particularly for a high-dimensional portfolio. We therefore offer a mathematics explanation to thoroughly investigate the performance of plug-in portfolios for a large p . Such an asymptotic analysis has been conducted in [18] for the extreme case of $p/n \rightarrow \infty$.

3.2. Known Σ , unknown β . In this case, we replace $\hat{\Sigma}_n$ by the background covariance matrix Σ in (3.3)-(3.5). Then we study the asymptotic performance of the plug-in portfolios when the number of assets (p) increases.

Theorem 3.1. If $\Theta_p = o(\frac{p}{n})$ and $\hat{\Sigma}_n = \Sigma$, we have as $p/n \rightarrow \infty$,

1. $\widehat{V}_{\log}(t, x) \xrightarrow{\mathbb{P}} -\infty$,
2. $\widehat{V}_{\text{pow}}(t, x) \xrightarrow{\mathbb{P}} \begin{cases} 0 & \text{for } 0 < \gamma < 1, \\ -\infty & \text{for } \gamma > 1, \end{cases}$
3. $\widehat{V}_{\text{exp}}(t, x) \xrightarrow{\mathbb{P}} -\infty$.

Proof. Notice that $\hat{\beta}_n \sim N(\beta, \frac{1}{n\delta t}\Sigma)$. Let ζ_j be the j th entry of $\Sigma^{-\frac{1}{2}}\beta$. Then $\Theta_p = \sum_{j=1}^p \zeta_j^2$ and the j th entry of $\Sigma^{-\frac{1}{2}}\hat{\beta}_n$ can be written as $\zeta_j + (n\delta t)^{-1/2}\epsilon_j$, where $\{\epsilon_j\}_{j=1}^p$ are independent standard normal random variables. Therefore,

$$\begin{aligned} \beta'\Sigma^{-1}\hat{\beta}_n &= \sum_{j=1}^p (\zeta_j^2 + (n\delta t)^{-1/2}\zeta_j\epsilon_j) = \Theta_p + \mathcal{O}_{\mathbb{P}}\left(\sqrt{\frac{\Theta_p}{n}}\right), \\ \hat{\beta}_n'\Sigma^{-1}\hat{\beta}_n &= \sum_{j=1}^p (\zeta_j + (n\delta t)^{-1/2}\epsilon_j)^2 = \Theta_p + \frac{1}{n\delta t} \sum_{j=1}^p \epsilon_j^2 + \frac{1}{\sqrt{n\delta t}} \sum_{j=1}^p \zeta_j\epsilon_j \\ &= \Theta_p + \frac{p}{n\delta t}[1 + o_{\mathbb{P}}(1)]. \end{aligned}$$

The quantities of our interest can be expressed as

$$(3.6) \quad \beta'\Sigma^{-1}\hat{\beta}_n - \frac{\hat{\beta}_n'\Sigma^{-1}\hat{\beta}_n}{2} = \frac{\Theta_p}{2} - \frac{p}{2n\delta t}[1 + o_{\mathbb{P}}(1)] = -\frac{p}{2n\delta t} + o_{\mathbb{P}}(1).$$

Then we have as $p/n \rightarrow \infty$, $\beta'\Sigma^{-1}\hat{\beta}_n - \frac{\hat{\beta}_n'\Sigma^{-1}\hat{\beta}_n}{2} \xrightarrow{\mathbb{P}} -\infty$. Direct application of the continuous mapping theorem to (3.3)–(3.5) yields the desired results. ■

Straightforward corollaries under the conditions of Theorem 3.1 are

$$\mathbb{E}[\widehat{V}_{\log}(t, x)] \rightarrow -\infty \text{ and } \mathbb{E}[\widehat{V}_{\text{pow}}(t, x)] \rightarrow -\infty$$

for $\gamma > 1$ as $p/n \rightarrow \infty$. This is consistent with the results proven in [24] for log utility and [18] for power utility. Notice that the sufficient condition in Theorem 3.1 is weaker than those in [18, 24].

Theorem 3.1 asserts that the plug-in optimal EUM portfolio tends to be the worst, in the sense of objective function (conditional expected utility) for $p/n \rightarrow \infty$. We noted in [18, 24] that the loss in expected utility is linear for log utility and exponential for power utility as the number of assets grows. However, there is hardly any improvement in the estimation of the mean vector, as only a marginal estimate is required. The empirical study in [16] shows that using the James–Stein estimator [28] for the mean does not improve the empirical performance of the optimal portfolio. Our empirical study draws the same conclusion. Many researchers thus consider trading strategies that are robust to the mean estimate, as in [27] and [23].

3.3. Known β , unknown Σ . We now turn to the problems in the plug-in portfolios, due to the accumulated estimation errors of the high-dimensional covariance matrix. The number of parameters in Σ increases quadratically as p increases and contributes to the instability of the sample estimate under the high-dimensional setting. When $p > n$, (3.2) is singular and we must consider its generalized inverse or Moore–Penrose pseudoinverse ($\hat{\Sigma}_n^-$), which is a biased

estimator, for estimating the precision matrix $\Omega = \Sigma^{-1}$. The unexpected outcomes produced by the bias are often studied in the statistical literature; see [4]. We assume β is known (i.e., $\hat{\beta}_n = \beta$ in (3.3)–(3.5)) to investigate the marginal effect of the Σ estimate.

To study the random matrix theory associated with our problem, the following lemma is useful.

Lemma 3.2.

1. $\beta' \hat{\Sigma}_n^- \Sigma \hat{\Sigma}_n^- \beta / \beta' \hat{\Sigma}_n^- \beta \xrightarrow{\mathbb{P}} 0$ as $p/n \rightarrow \infty$.
2. If $\Theta_p = o(\frac{p}{n})$, then $\beta' \hat{\Sigma}_n^- \beta / \sqrt{\beta' \hat{\Sigma}_n^- \Sigma \hat{\Sigma}_n^- \beta} \xrightarrow{\mathbb{P}} 0$ (and in $\mathcal{L}^2$) as $p/n \rightarrow \infty$.
3. If $\Theta_p = o(\frac{p}{n})$, then $\sqrt{\beta' \hat{\Sigma}_n^- \Sigma \hat{\Sigma}_n^- \beta} \xrightarrow{P} 0$ and $\beta' \hat{\Sigma}_n^- \beta \xrightarrow{\mathbb{P}} 0$ as $p/n \rightarrow \infty$.

Proof. The third assertion is a direct consequence of the first and second assertions according to Slutsky's theorem. Hence, we only prove the first and second assertions.

We first consider a simpler case that the covariance matrix is an identity matrix: $\Sigma = I$. The eigenvectors of $\hat{\Sigma}_n$, denoted as $\hat{\xi}_1, \dots, \hat{\xi}_n$, which correspond to nonzero eigenvalues $\hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_n$ of $\hat{\Sigma}_n$, span a subspace of $\mathbb{R}^p$. Moreover, $(\hat{\lambda}_1, \dots, \hat{\lambda}_n)$ and $(\hat{\xi}_1, \dots, \hat{\xi}_n)$ are independent, while $\hat{\xi}_j$ are identically distributed uniformly on the unit p -sphere. The Moore–Penrose pseudoinverse of $\hat{\Sigma}_n$ can then be expressed as $\hat{\Sigma}_n^- = \sum_{i=1}^n \frac{1}{\hat{\lambda}_i} \hat{\xi}_i \hat{\xi}_i'$.

1. As proven in the Theorem 1(ii) in [31], $\hat{\lambda}_i = \mathcal{O}_P(\frac{p}{n})$ for all i as $\frac{p}{n} \rightarrow \infty$. We recognize that

$$\frac{\beta' \hat{\Sigma}_n^- \hat{\Sigma}_n^- \beta}{\beta' \hat{\Sigma}_n^- \beta} = \frac{\sum_{i=1}^n \frac{1}{\hat{\lambda}_i^2} (\beta' \hat{\xi}_i)^2}{\sum_{i=1}^n \frac{1}{\hat{\lambda}_i} (\beta' \hat{\xi}_i)^2} \leq \frac{\hat{\lambda}_1}{\hat{\lambda}_n^2} = \mathcal{O}_P\left(\frac{n}{p}\right).$$

Hence, $\beta' \hat{\Sigma}_n^- \hat{\Sigma}_n^- \beta / \beta' \hat{\Sigma}_n^- \beta$ converges to 0 in probability as $p/n \rightarrow \infty$.

2. According to the Cauchy–Schwarz inequality,

$$\left[\frac{\beta' \hat{\Sigma}_n^- \beta}{\sqrt{\beta' \hat{\Sigma}_n^- \hat{\Sigma}_n^- \beta}} \right]^2 = \frac{[\sum_{i=1}^n \frac{1}{\hat{\lambda}_i} (\beta' \hat{\xi}_i)^2]^2}{\sum_{i=1}^n \frac{1}{\hat{\lambda}_i^2} (\beta' \hat{\xi}_i)^2} \leq \sum_{i=1}^n (\beta' \hat{\xi}_i)^2.$$

The expectation of the right-hand side in the above inequality is equal to $(n/p)\|\beta\|^2$.

Hence, $\beta' \hat{\eta}_n / \sqrt{\hat{\eta}_n' \hat{\eta}_n}$ converges to 0 in $\mathcal{L}^2$ once $\|\beta\|^2 = o(\frac{p}{n})$.

Yet, consider the general positive definite Σ . We take $\beta^* = \Sigma^{-\frac{1}{2}} \beta$ and $\hat{\Sigma}_n^* = \Sigma^{-\frac{1}{2}} \hat{\Sigma}_n \Sigma^{-\frac{1}{2}}$, where $\hat{\Sigma}_n^*$ is a Gram matrix of pseudosamples $\{r^{*(l)}\}_{l=1}^n$ and $r^{*(l)} = (\delta t \Sigma)^{-\frac{1}{2}} r^{(l)}$ has the identity covariance matrix. Notice that $\frac{\beta' \hat{\Sigma}_n^- \Sigma \hat{\Sigma}_n^- \beta}{\beta' \hat{\Sigma}_n^- \beta} = \frac{\beta^{*'} \hat{\Sigma}_n^{*-} \hat{\Sigma}_n^{*-} \beta^*}{\beta^{*'} \hat{\Sigma}_n^{*-} \beta^*}$ and $\frac{\beta' \hat{\Sigma}_n^- \beta}{\sqrt{\beta' \hat{\Sigma}_n^- \Sigma \hat{\Sigma}_n^- \beta}} = \frac{\beta^{*'} \hat{\Sigma}_n^{*-} \beta^*}{\sqrt{\beta^{*'} \hat{\Sigma}_n^{*-} \hat{\Sigma}_n^{*-} \beta^*}}$. According to the results of the case of the identity covariance matrix, if the condition for $\|\beta^*\|^2 = \Theta_p$ is satisfied, then $\beta' \hat{\Sigma}_n^- \Sigma \hat{\Sigma}_n^- \beta / \beta' \hat{\Sigma}_n^- \beta$ and $\beta' \hat{\Sigma}_n^- \beta / \sqrt{\beta' \hat{\Sigma}_n^- \Sigma \hat{\Sigma}_n^- \beta}$ converges to 0 in probability as $p/n \rightarrow \infty$. ■

Lemma 3.2 and the continuous mapping theorem result in the following theorem about the performance of the plug-in portfolios subject to the accumulated estimation errors of $\hat{\Sigma}_n^-$ in p -asymptotics.

Theorem 3.3. If $\Theta_p = o(\frac{p}{n})$ and $\hat{\beta}_n = \beta$, we have as $p/n \rightarrow \infty$,

1. $\widehat{V}_{\log}(t, x) \xrightarrow{\mathbb{P}} \log x + r_f(T - t),$
2. $\widehat{V}_{\text{pow}}(t, x) \xrightarrow{\mathbb{P}} (xe^{r_f(T-t)})^{1-\gamma}/(1-\gamma),$
3. $\widehat{V}_{\text{exp}}(t, x) \xrightarrow{\mathbb{P}} -\frac{1}{\gamma} \exp \left\{ -\gamma xe^{r_f(T-t)} \right\}.$

The limiting functions in Theorem 3.3 are the value functions, with all money going into the bank account. Here, all the plug-in optimal portfolios suggest no investment in risky assets under the estimation errors of the covariance matrix. In terms of risk exposure, the effect of the estimation in the mean is more serious than that in the covariance matrix, as the latter is completely risk-free.

It has been proven in [35] that the accumulated estimation errors of the covariance matrix make the optimal mean-variance portfolio degenerate to a random stock picking strategy. Interestingly, its effect on the Merton problem is rather different, which deserves separate analysis.

When both β and Σ are unknown, intuitively the asymptotic performance of the plug-in portfolios cannot be better than the marginal case in either Theorem 3.1 or Theorem 3.3. Hence, the plug-in optimal portfolio should be used in the worst case. Both the sample estimates of β and Σ can be seen as overestimations. However, the accumulated effect of these estimations poses an analytical challenge, which we examine in our numerical studies.

Although the effect of improved estimators, such as the Stein-type biased estimators, is not analyzed in this section, the James–Stein estimator for the drifts has shown no empirical improvement in [16] and we draw the same conclusion for the Merton problem in a later section. One may also consider Stein-type shrinkage estimator for the correlation matrix. However, it has been detailed in [25] that the computational cost could be huge and the shrinking method has to be chosen carefully among many possibilities for a specific purpose. To the best of our knowledge, there is no shrinkage estimator for the correlation matrix that is shown to have the oracle property for dynamic EUM in the literature. In other words, it is not clear if Stein-type shrinkage estimate of the correlation matrix helps recovering the true background optimal strategy asymptotically.

We admit that although the number of risky assets p could be large, p/n may be far from infinity in practice.¹ In the following section, we propose a novel approach which asymptotically recovers the background true optimal policy which is free of estimation error for all cases of $\frac{p}{n} \rightarrow 0$, a constant, or infinity as long as $\frac{\log p}{n} \rightarrow 0$.

4. Resolution: Constrained ℓ_1 -minimization. In practice, traditional estimation methods often fail, as described in section 3 and in many empirical studies. To remedy this, we take into account an important observation: the effective parameters in the optimal portfolios (2.4) are $\Sigma^{-1}\beta$. We view this as a whole unknown vector η . Instead of separate estimation, we aim to directly control the estimation errors of η , which clearly reduces the number of parameters to be estimated from $p(p+3)/2$ (in Ω and β) to p (in η and β) and avoids the computation of an inverse of a large (singular) matrix. We denote by $\hat{\eta}$ and $\hat{\beta}'\hat{\eta}$ the estimates of η and Θ_p , respectively, which are substituted into the optimal portfolios for implementation.

We adopted the constrained ℓ_1 -minimization (or LPO) approach, proposed in [35], to directly estimate $\eta = \Sigma^{-1}\beta$:

¹We thank an anonymous referee of pointing this out.

$$(4.1) \quad \tilde{\eta} \in \arg \min_{\eta \in \mathbb{R}^p} \left\{ |\eta|_1 \text{ subject to } |\hat{\Sigma}_n \eta - \hat{\beta}_n|_\infty \leq \lambda_n \right\},$$

where $\hat{\beta}_n$ and $\hat{\Sigma}_n$ are defined in (3.1) and (3.2), λ_n is a tuning parameter, $|a|_1 = \sum_{i=1}^p |a_i|$, and $|a|_\infty = \sup_i |a_i|$ for $a \in \mathbb{R}^p$. This approach still requires the sample-based estimates of β and Σ , but it is not limited to these and any other improved estimators could be substituted.

The ℓ_1 -norm appears in (4.1), but we do not amend the Merton problem by imposing norm constraints, as in [15, 18, 20]. We instead construct a novel direct estimation method with oracle properties, which can be applied to various optimal portfolios derived from the optimization theory in the literature. The extension of this approach to robust portfolios is discussed in Appendix A. The proposed approach essentially renders a data-driven portfolio policy under a particular structure. The tuning parameter λ_n is determined by cross-validation (CV) in practice; hence the proposed approach is execution-oriented and based on the current information. We use $\tilde{\eta}$ as the LPO estimator, as its computation can be realized with a linear programming procedure, given in section 5.1. The resulting optimal EUM portfolios (2.4) are then referred to as the LPO portfolios. Using the notation of $\tilde{\eta}$, the LPO portfolios read

$$(4.2) \quad \begin{aligned} \widetilde{u^*(t, x)} &= \widetilde{R(t, x)\tilde{\eta}}, \text{ where } \widetilde{R(t, x)} \text{ solves} \\ \widetilde{R_t} + \frac{\hat{\beta}'_n \tilde{\eta}}{2} \widetilde{R^2 R_{xx}} + r_f x \widetilde{R_x} - r_f \widetilde{R} &= 0, \quad \widetilde{R(T, x)} = -U'(x)/U''(x), \\ \text{and specifically } \widetilde{u_{log}^*(t, x)} &= x \tilde{\eta}, \quad \widetilde{u_{pow}^*(t, x)} = \frac{x}{\gamma} \tilde{\eta}, \quad \widetilde{u_{exp}^*(t, x)} = \frac{e^{-r_f(T-t)}}{\gamma} \tilde{\eta}. \end{aligned}$$

Their corresponding performance measures are given by (3.3)–(3.5) with $\hat{\Sigma}_n^{-1} \hat{\beta}_n$ replaced by $\tilde{\eta}$ and are denoted by $\widetilde{V_{log}}$, $\widetilde{V_{pow}}$, and $\widetilde{V_{exp}}$, respectively.

The LPO approach finds the η estimate that best approximates the background one, with the least cost in terms of the ℓ_1 -norm. The optimal EUM portfolios (2.4) are all proportional to η , so the resulting LPO portfolios exhibit sparsity, where the cardinality agrees with that of the η estimate. A natural concern is the comparison between the LPO portfolio and the ℓ_1 -norm-constrained portfolios, discussed in [18, 24], as both solutions generate sparse optimal policies.

The ℓ_1 -norm-constrained portfolios are essentially LASSO-type problems requiring the inverse of the covariance matrix estimate (Cholesky decomposition of $\hat{\Sigma}_n$) as an input. If, due to estimation errors, the inverse does not exist, the generalized inverse is indispensable in the implementation phase. In spite of the expensive computational cost of the generalized matrix inverse, our numerical examples show that it fails to work for a large p . The LPO approach, inspired by the Dantzig selector [13], is slightly more sparse than the norm-constrained portfolio. This eventually dilutes the effect associated with the estimation errors in the mean vector more effectively. More importantly, it does not require the $\hat{\Sigma}_n$ to be invertible and guarantees the existence of the LPO estimator. Being an enhanced estimator, the LPO approach still needs the (sample) estimates of Σ and β as inputs. Therefore, the sample size should be reasonably large ($n \geq \log p$) to rule out data deficiency. We refer readers to [10, 11, 13] for the application of the ℓ_1 -minimization methods in other statistics disciplines.

When $\lambda_n > |\hat{\beta}_n|_\infty$, the optimal solution becomes $\tilde{\eta} = 0$, which infers a noninvestment strategy. The cardinality of $\tilde{\eta}$ increases as λ_n decreases. The lower boundary of the feature

region for λ_n ($\lambda_{\min}$) is determined by decreasing the λ_n from $|\hat{\beta}_n|_\infty$ until it reaches zero, or until the corresponding feasible set for the problem (4.1) is empty. Consequently, we are interested in $\lambda_n \in [\lambda_{\min}, |\hat{\beta}_n|_\infty]$. If $\lambda_{\min} = 0$, then the problem (4.1) with $\lambda_n = 0$ has a feasible solution in the separate estimation method: $\tilde{\eta} = \hat{\Sigma}_n^{-1} \hat{\beta}_n$. Therefore, the LPO approach nests the separate estimate as a solution candidate when we empirically cross-validate the λ_n . However, $\lambda_{\min}$ is generally nonzero for a high-dimensional portfolio and particularly for $p > n$, due to the possible singularity of $\hat{\Sigma}_n$. It turns out that the LPO approach will only select k assets ($k < p$) into the optimal portfolio, resulting in a genuine portfolio selection.

4.1. Oracle properties. We study the oracle properties of the LPO approach to the investment problems in section 2.2. The oracle benchmarks are defined as the optimal value functions with known parameters, specified in (2.6), (2.8), and (2.10). Notice that the historical daily returns D under the model (2.1) follow the normal distribution with the mean $\delta t(\mu - \text{diag}(\Sigma)/2)$ and the covariance matrix $\delta t\Sigma$. It is clear that the daily return vectors satisfy the moment condition (M1) in [35]:

$$(M1) \quad \log p \leq n, \quad \mathbb{E} \exp(c_1 \bar{\Psi}^2) \leq K_1, \quad \text{and} \quad \mathbb{E} \exp(c_1 \Psi_i^2 / \sigma_{ii}^0) \leq K_1, \quad \forall i, \\ \text{for some constants } c_1, K_1 > 0,$$

where in our case $\Psi_i = \sigma_i W_1$, $\Psi = (\Psi_1, \dots, \Psi_p)' = \sigma W_1$, $\bar{\Psi} = \Psi' \Omega \beta / \sqrt{\Theta_p}$ and we read $\Sigma = \{\sigma_{ij}^0\}_{1 \leq i, j \leq p}$. To examine under a high-dimensional setting, we introduce a condition which is also imposed in [10, 11, 35]:

$$(C1) \quad \log p \leq n, \quad \max_{1 \leq i \leq p} \sigma_{ii}^0 \leq M \quad \text{and} \quad \Theta_p \geq K \quad \text{for some constants } M, K > 0.$$

The following lemma is crucial in the derivation of the oracle properties of the LPO approach.

Lemma 4.1. *Suppose that the historical daily return data satisfies (2.1) and condition (C1) holds. With probability greater than $1 - \mathcal{O}(p^{-1})$, we have*

$$(4.3) \quad |\hat{\Sigma}_n \eta - \hat{\beta}_n|_\infty \leq C \sqrt{\Theta_p \log p / n}$$

for a sufficiently large constant C , where $\hat{\beta}_n$ and $\hat{\Sigma}_n$ are defined in (3.1) and (3.2), and $\eta = \Omega \beta$ is the background effective parametric vector.

Remark. Hereafter, C represents a large constant and can be varied from place to place.

Proof. This is similar to the Lemma 2 in [11] with $n_1 = n$, $Y_l = 0$, $X_l = r^{(l)} - r_f \mathbf{1}_p = \delta t(\beta - \frac{1}{2} \text{diag}(\Sigma)) + \sqrt{\delta t} \Psi^{(l)}$, $l = 1, \dots, p$. We notice that

$$\hat{\Sigma}_n = \frac{1}{n} \sum_{l=1}^n \Psi^{(l)} \Psi^{(l)'} - \tilde{\Psi} \tilde{\Psi}' =: \tilde{\Sigma}_n - \tilde{\Psi} \tilde{\Psi}', \quad \text{where} \quad \tilde{\Psi} = \frac{1}{n} \sum_{l=1}^n \Psi^{(l)}, \\ \hat{\beta}_n = \beta + \frac{1}{2} \text{diag}(\hat{\Sigma}_n - \Sigma) + \frac{1}{\sqrt{\delta t}} \tilde{\Psi}.$$

Then we have

$$\begin{aligned} |\hat{\Sigma}_n \eta - \hat{\beta}_n|_\infty &= |\tilde{\Sigma}_n \eta - \beta - \tilde{\Psi} \tilde{\Psi}' \eta - \hat{\beta}_n + \beta|_\infty \\ &\leq |\tilde{\Sigma}_n \eta - \beta|_\infty + |\tilde{\Psi} \tilde{\Psi}' \eta - \frac{1}{\sqrt{\delta t}} \tilde{\Psi}|_\infty + \frac{1}{2} |\hat{\Sigma}_n - \Sigma|_\infty. \end{aligned}$$

The first two terms are handled similarly to [11]. To control the last term, we use a lemma proven in [10]: under condition (M1) with a probability greater than $1 - \mathcal{O}(p^{-1})$, we have

$$|\hat{\Sigma}_n - \Sigma|_\infty \leq C \sqrt{\log p/n}.$$

Then the result follows. ■

This lemma implies that if we set $\lambda_n = C \sqrt{\Theta_p \log p/n}$, then the background η is a candidate in the feasible set of (4.1), which is useful to prove the oracle properties. The idea of proving oracle properties is to show the convergence of the featured quantities in the out-of-sample measures with the LPO estimators.

Theorem 4.2. *Let $\lambda_n = C \sqrt{\Theta_p \log p/n}$ with $C > 0$ being a sufficiently large constant, $a_n := |\hat{\Sigma}_n - \Sigma|_\infty$, and $b_n := |\hat{\beta}_n - \beta|_\infty$. Suppose that condition (C1) holds and*

$$(4.4) \quad d_n := (\lambda_n + b_n) |\eta|_1 + a_n |\eta|_1^2 = o_{\mathbb{P}}(1).$$

Then we have

1. $\widehat{V}_{\log}(t, x) = \widehat{V}_{\log|_{\{\hat{\Sigma}_n^{-1} \hat{\beta}_n = \tilde{\eta}\}}}(t, x) = V_{\log}(t, x) - \mathcal{O}_{\mathbb{P}}(d_n),$
2. $\widehat{V}_{\text{pow}}(t, x) = \widehat{V}_{\text{pow}|_{\{\hat{\Sigma}_n^{-1} \hat{\beta}_n = \tilde{\eta}\}}}(t, x) = \begin{cases} V_{\text{pow}}(t, x)(1 - \mathcal{O}_{\mathbb{P}}(d_n)) & \text{for } 0 < \gamma < 1, \\ V_{\text{pow}}(t, x)(1 + \mathcal{O}_{\mathbb{P}}(d_n)) & \text{for } \gamma > 1, \end{cases}$
3. $\widehat{V}_{\text{exp}}(t, x) = \widehat{V}_{\text{exp}|_{\{\hat{\Sigma}_n^{-1} \hat{\beta}_n = \tilde{\eta}\}}}(t, x) = V_{\text{exp}}(t, x)(1 + \mathcal{O}_{\mathbb{P}}(d_n)).$

Proof. Using (4.3), we have

$$|\eta' \hat{\Sigma}_n \tilde{\eta} - \beta' \tilde{\eta}| \leq (\lambda_n + |\hat{\beta}_n - \beta|_\infty) |\tilde{\eta}|_1 \leq (\lambda_n + b_n) |\eta|_1.$$

By the definition of $\tilde{\eta}$, we have

$$|\eta' \hat{\Sigma}_n \tilde{\eta} - \eta' \beta| \leq (\lambda_n + |\hat{\beta}_n - \beta|_\infty) |\eta|_1 \leq (\lambda_n + b_n) |\eta|_1.$$

Hence, combining the above two inequalities, we have

$$(4.5) \quad |\beta' \tilde{\eta} - \Theta_p| \leq 2(\lambda_n + b_n) |\eta|_1 \leq C(\lambda_n + b_n) |\eta|_1.$$

Next, we notice that

$$|\Sigma \tilde{\eta} - \hat{\beta}_n|_\infty \leq |\Sigma - \hat{\Sigma}_n|_\infty |\tilde{\eta}|_1 + \lambda_n \leq a_n |\eta|_1 + \lambda_n,$$

and thus

$$(4.6) \quad \begin{aligned} |\tilde{\eta}' \Sigma \tilde{\eta} - \hat{\beta}_n' \tilde{\eta}| &\leq a_n |\eta|_1^2 + \lambda_n |\eta|_1 \leq d_n, \\ |\tilde{\eta}' \Sigma \tilde{\eta} - \Theta_p| &\leq a_n |\eta|_1^2 + (3\lambda_n + b_n) |\eta|_1 \leq C d_n. \end{aligned}$$

Now, we investigate the asymptotic performance of the LPO portfolios:

$$\begin{aligned}
|\widetilde{V}_{log} - V_{pow}| &= (T-t) \left| \beta' \tilde{\eta} - \Theta_p - \frac{1}{2} (\tilde{\eta}' \Sigma \tilde{\eta} - \Theta_p) \right| \\
&\leq (T-t) \left[|\beta' \tilde{\eta} - \Theta_p| + \frac{1}{2} |\tilde{\eta}' \Sigma \tilde{\eta} - \Theta_p| \right] = \mathcal{O}_{\mathbb{P}}(d_n), \\
\left| \frac{\widetilde{V}_{pow}}{V_{pow}} - 1 \right| &= \left| \exp \left\{ \frac{(1-\gamma)(T-t)}{\gamma} \left[\beta' \tilde{\eta} - \Theta_p - \frac{1}{2} (\tilde{\eta}' \Sigma \tilde{\eta} - \Theta_p) \right] \right\} - 1 \right| \\
&\leq \mathcal{O}_{\mathbb{P}} \left(\frac{(1-\gamma)T}{\gamma} \left[|\beta' \tilde{\eta} - \Theta_p| + \frac{1}{2} |\tilde{\eta}' \Sigma \tilde{\eta} - \Theta_p| \right] \right) = \mathcal{O}_{\mathbb{P}}(d_n), \\
\left| \frac{\widetilde{V}_{exp}}{V_{exp}} - 1 \right| &= \left| \exp \left\{ -(T-t) \left[\beta' \tilde{\eta} - \Theta_p - \frac{1}{2} (\tilde{\eta}' \Sigma \tilde{\eta} - \Theta_p) \right] \right\} - 1 \right| \\
&\leq \mathcal{O}_{\mathbb{P}} \left((T-t) \left[|\beta' \tilde{\eta} - \Theta_p| + \frac{1}{2} |\tilde{\eta}' \Sigma \tilde{\eta} - \Theta_p| \right] \right) = \mathcal{O}_{\mathbb{P}}(d_n),
\end{aligned}$$

where (4.5) and (4.6) are used. Then the results follow. ■

The oracle properties assert that the LPO portfolios enable the corresponding out-of-sample value functions to approach their true values for a small d_n . In fact, the errors of β and Σ estimates are normally of order $\sqrt{\log p/n}$, and thus the rate of convergence d_n is minimal when $\sqrt{\log p/n}$ is small. We can expect the LPO portfolios perform as well as the theoretically optimal portfolios when $\log p/n$ is small. Together with the results in section 3, Theorem 4.2 shows that the LPO portfolios have dominating performance over the traditional plug-in portfolios when p/n is large and $\log p/n$ is small. In the later section on numerical and empirical studies, we will also show that the LPO portfolios are superior in the small p case too.

The magnitude of d_n depends on the precision of the estimations in Σ and β and the cardinality of η . If the estimation contains minimal error, then a large cardinality for η is allowed. Alternatively, large estimation errors in Σ and β can be compensated for by shrinking the cardinality of η . The proof is independent of the choice of estimators for Σ and β . Therefore, the oracle properties hold true for any estimators of Σ and β that satisfy the inequality (4.3). In fact, the LPO approach is general enough to include the separate estimation and the shrinkage method in [30] as special cases.

If we replace condition (C1) by

$$\begin{aligned}
\text{(C2)} \quad & \log p \leq n, \quad M^{-1} \leq \lambda_{\min}(\Sigma) \leq \lambda_{\max}(\Sigma) \leq M \text{ and } \Theta_p \geq K, \\
& \text{for some constants } M, K > 0,
\end{aligned}$$

where $\lambda_{\min}(\Sigma)$ and $\lambda_{\max}(\Sigma)$ are the smallest and largest eigenvalues of Σ , respectively, Theorem 4.2 is still true and we have $|\eta|_1^2 \leq M^3 |\eta|_0 \Theta_p$, where $|\eta|_0$ is the cardinality of η . In other words, the larger the Θ_p the slower the convergence rate of the LPO estimates. This is reasonable, because if the oracle portfolios perform extraordinarily well it is difficult for any data-driven method to achieve a similarly high standard of performance.

It is noteworthy from the sufficient condition (4.4) that if the estimation errors in the mean vector b_n are large, additional sparsity of the optimal EUM portfolio is required, to ensure convergence. The nature of the LPO approach is to effectively control the utility loss due to the mean estimation errors.

5. Numerical studies. To compare continuous-time strategies under different implementation and/or estimation schemes, we rebalance the portfolios daily, though this suffers from mild but unavoidable discretization errors. Based on out-of-sample statistics, we examine the finite-sample performance of our LPO approach, verify our theoretical results, and compare them with alternative approaches through simulation studies.

5.1. Implementation. We detail the implementation of the LPO estimator here. The LPO estimator of $\tilde{\eta}$ is the solution of the minimization problem:

$$\min |\eta|_1 \quad \text{subject to: } |\hat{\Sigma}_n \eta - \hat{\beta}_n|_\infty \leq \lambda_n.$$

By combining ℓ_1 and ℓ_∞ norms we attempt to apply linear programming, and our theoretical results can be generalized to the combination of ℓ_p and ℓ_q norms where $1/p + 1/q = 1$. The convenient linear programming formulation allows us to solve the problem using the parametric simplex method, which is efficient and guarantees machine precision with a primal-dual gap. More details can be found in [37]. To reformulate our problem into the parametric simplex problem in [37], let $\eta = \eta^+ - \eta^-$, where $\eta^\pm = \max(\pm\eta, 0)$ and one of η^+, η^- be 0. The target minimization problem can be rewritten as

$$\max (-\mathbf{1}_p \quad -\mathbf{1}_p) \begin{pmatrix} \eta^+ \\ \eta^- \end{pmatrix} \quad \text{subject to} \quad \begin{pmatrix} \hat{\Sigma}_n & -\hat{\Sigma}_n \\ -\hat{\Sigma}_n & \hat{\Sigma}_n \end{pmatrix} \begin{pmatrix} \eta^+ \\ \eta^- \end{pmatrix} \leq \begin{pmatrix} \lambda_n \mathbf{1}_p + \hat{\beta}_n \\ \lambda_n \mathbf{1}_p - \hat{\beta}_n \end{pmatrix}.$$

For a given λ_n , the parametric simplex method computes the optimal (η^+, η^-) so that we recover the optimal $\tilde{\eta}$ as $\tilde{\eta} = \eta^+ - \eta^-$.

As discussed in section 4, the feasible set of λ_n is $\Lambda := [\lambda_{\min}, |\hat{\beta}_n|_\infty]$, where $\lambda_{\min}$ is determined through the trials using historical data. We select the optimal tuning parameter λ_n among Λ by CV, which is a standard approach in machine learning. Specifically, we order the data $D = \{r^{(1)}, \dots, r^{(n)}\}$ into f groups, where f is the number of folds of CV, as the data are ordered in time. For each iteration, the groups with sizes n/f and $n - n/f$ act as test data and training data, respectively. For each λ_n , we run our program using the training data and then examine the corresponding LPO portfolios $(u_{\lambda_n}^*(t, x))$ using the test data of sample size n/f . The optimal λ_n has the portfolio that optimizes the investor's objective across all folds of the test data.

For a given utility function, we denote the objective function value at the k th iteration as $\widehat{\text{Obj}}_k(\lambda_n)$. Specifically, $\widehat{\text{Obj}}_k(\lambda_n) = \widehat{U}(X_k^{u_{\lambda_n}^*}(T))$, where $X_k^{u_{\lambda_n}^*}(T)$ is the terminal wealth found by adopting the LPO portfolio $u_{\lambda_n}^*(t, x)$, computed using the test data. For example, $U(x) = -e^{-\gamma x}/\gamma$ for an exponential utility and $U(x) = x^{1-\gamma}/(1-\gamma)$ for a power utility. To match the investment horizon, f should be set as $n/n.test$, where $n.test$ is the number of days within the investment horizon. We suffer from a large computational burden as $n/n.test$

is usually large, so we instead aim to find an optimal LPO estimator $\tilde{\eta}(\lambda^*)$ that maximizes (3.3)–(3.5) or equivalently maximizes

$$\beta' \tilde{\eta}(\lambda_n) - \tilde{\eta}(\lambda_n)' \Sigma \tilde{\eta}(\lambda_n) / 2$$

with respect to λ_n . Hence, we use training data with size $n - n/f$ to compute the LPO estimate $\tilde{\eta}(\lambda_n)$ and use test data with size n/f to compute the proxies (sample estimates) of β and Σ : $\hat{\beta}_{test}$ and $\hat{\Sigma}_{test}$. In this study, we choose $f = 5$. Specifically, we take

$$\widehat{\text{Obj}}_k(\lambda_n) = \hat{\beta}'_{test} \tilde{\eta}(\lambda_n) - \tilde{\eta}(\lambda_n)' \hat{\Sigma}_{test} \tilde{\eta}(\lambda_n) / 2.$$

Then the CV choice of λ_n is given by

$$\lambda_{CV} = \max_{\lambda_n} \sum_{k=1}^f \widehat{\text{Obj}}_k(\lambda_n).$$

If the maximum is attained at several λ s, we select the smallest.

A distinct advantage of the LPO approach is its applicability to the general utility function. Given the LPO estimate of Θ_p : $\widetilde{\Theta}_p = \hat{\beta}'_n \tilde{\eta}$, $\tilde{R}(t, x)$ can be solved numerically from the PDE (4.2), using the implicit finite difference method; see [6]. However, the utility function must be specified to offer the boundary conditions. For instance, we introduce a family of utility functions that allows for nonlinear risk tolerance (nonconstant relative risk aversion): mixture of power utilities. They are specified as follows.

$$(5.1) \quad U(x) = c_1 \frac{x^{1-\gamma_1}}{1-\gamma_1} + c_2 \frac{x^{1-\gamma_2}}{1-\gamma_2}, \quad c_1, c_2 > 0, \quad \gamma_1 \geq \gamma_2 > 0, \quad \gamma_1, \gamma_2 \neq 1.$$

This family of utilities has recently been of great interest [9]. Quantitative treatments of related problems can be found in [22, 34]. Due to the analytical intractability, the norm-constrained solution associated with the mixture of power utilities is not available in the literature, but the application of the LPO approach to this class of utility functions is straightforward. A typical mixture of power utilities combines an unbounded above-positive utility ($\gamma_2 < 1$) and an unbounded below-negative utility ($\gamma_1 > 1$), producing a nonlinear risk aversion. In such a situation, the risk-tolerance function at time T is

$$R(T, x) = -\frac{U'(x)}{U''(x)} = \frac{c_1 x^{-(\gamma_1-\gamma_2)} + c_2}{c_1 \gamma_1 x^{-(\gamma_1-\gamma_2)} + c_2 \gamma_2} x \sim \begin{cases} x/\gamma_2 & \text{as } x \rightarrow \infty, \\ x/\gamma_1 & \text{as } x \rightarrow 0, \end{cases}$$

from which the boundary conditions for the numerical scheme are deduced as

$$R(t, 0) = 0, \quad R(t, x^*) = x^*/\gamma_2, \quad R(T, x) = \frac{c_1 x^{-(\gamma_1-\gamma_2)} + c_2}{c_1 \gamma_1 x^{-(\gamma_1-\gamma_2)} + c_2 \gamma_2} x$$

for a sufficiently large x^* .

5.2. Simulation design. Numerical studies are constructed to compare the LPO portfolio with the oracle portfolio, and the plug-in portfolio with the ℓ_1 norm-constrained portfolio in [18, 24] for the power utility and the mixture of power utilities. The results for the log and exponential utilities are similar, though not reported here. Three correlation structures among the risky asset returns are used to accommodate sparse, nonsparse, and even random correlation matrices. This choice of correlation matrices is similar to those of [11, 35].

The numerical experiments are designed as follows. Using the discretized model (2.1) for p stocks, we simulate two-year daily prices ($n = 504$) as training data for estimation and two-week daily prices ($n_{\text{test}} = 10, T = 1/25$) as test data for the out-of-sample evaluation. The risk-free interest rate is $r_f = 0.01$. All the p stocks have the same initial price of $s_i = 40$. The excess mean return rates are $\beta = (0.3, \dots, 0.3, 0, \dots, 0)'$ with $c_0 = 10$ as the number of nonzeros, in spite of the value of p . The covariance matrix of the risky asset returns takes the form $\Sigma = 0.2^2 \Gamma$, where $\Gamma = (\gamma_{ij})_{p \times p}$ is the correlation matrix of the following three possibilities:

Model 1. $\gamma_{ij} = 0.8^{|i-j|}$ for $1 \leq i, j \leq p$.

Model 2. $\gamma_{ii} = 1$ for $1 \leq i \leq p$ and $\gamma_{ij} = 0.5$ for $i \neq j$.

Model 3. $\Gamma^{-1} = (B + \delta I)/(1 + \delta)$, where $B = (b_{ij})_{p \times p}$ with independent $b_{ij} = b_{ji} = 0.5 \times \text{Ber}(1, 0.2)$ for $1 \leq i \leq c_0, i < j \leq p$; $b_{ij} = b_{ji} = 0.5$ for $c_0 + 1 \leq i < j \leq p$; $b_{ii} = 1$ for $1 \leq i \leq p$. Here $\text{Ber}(1, 0.2)$ is a Bernoulli random variable that takes a value of 1 with a probability of 0.2 and 0 with a probability of 0.8, and $\delta = \max(-\lambda_{\min}(B), 0) + 0.05$ to ensure that Γ^{-1} is a positive definite matrix. The matrix is finally standardized to unit diagonals.

In Model 1, Σ is approximately sparse and its inverse Ω is a 3-sparse matrix. In Model 2, Ω and $\eta = \Omega\beta$ are approximately sparse. In Model 3, the first c_0 rows and columns of Ω are sparse but the rest of the matrix is not sparse. Therefore, Σ is not sparse and contains random elements.

The initial wealth of the investor is $x_0 = 1$. The risk aversion coefficient of CRRA investor is $\gamma_{\text{pow}} = 2$. For the investor with the mixture of power utilities, we assume $c_1 = c_2 = 0.5$, $\gamma_1 = 1.2$ and $\gamma_2 = 0.25$ in (5.1).

5.3. Numerical results. We consider different levels of p (100, 200, 400, 600) and conduct 100 simulations in each case. In each simulation, we compute the terminal utility values and report the mean and standard deviation from the 100 simulation runs. In the following tables, OC represents the oracle portfolio; LPO represents the LPO portfolio; NC represents the ℓ_1 -norm constrained portfolio of [18]; and PI represents the traditional plug-in portfolio.

The numerical results for the power utility and the mixture of powers utility are reported in Table 1. For the optimal policy of the power utility maximization, the portfolio wealth must theoretically stay nonnegative. However, the data-driven portfolio may go negative due to discretization errors and potential estimation errors. We record the number of portfolios experiencing negative wealth during the investment period and show this in square brackets, [no.]. A negative portfolio leads the terminal wealth to a complex number, so this situation is eliminated from the mean and standard deviation estimates. A large number of negative portfolio wealths indicates highly distorted statistics. The NC portfolios cannot be implemented for $n \ll p$ as the inverse of the covariance matrix does not exist, and the numerical generalized inverse computation fails to render a numerical value.

Table 1

Mean (standard deviation) of the terminal utility values over 100 simulations. Negative wealth is shown in brackets. N.A. stands for “not applicable” for an undefined value.

p	Model 1						
	Power utility				Mixture of power utilities		
	OC	LPO	NC	PI	OC	LPO	PI
100	-0.91 (0.23)	-0.98 (0.19)	-16.89 (157)	-2.71 [3] (5.63)	-1.64 (0.76)	-1.78 (0.51)	-1.89 [13] (1.16)
200	-0.92 (0.25)	-0.98 (0.15)	-39.02 [2] (337)	-365 [54] (2400)	-1.63 (0.80)	-1.78 (0.40)	-1.93 [5] (0.34)
400	-0.94 (0.27)	-1.00 (0.12)	-14.69 [6] (106)	N.A. [100] (N.A.)	-1.70 (0.89)	-1.85 (0.34)	N.A. [100] (N.A.)
600	-0.90 (0.32)	-1.00 (0.14)	XXXXXXXXXX	0 [99] (N.A.)	-1.51 (0.89)	-1.81 (0.41)	N.A. [100] (N.A.)
p	Model 2						
	Power utility				Mixture of power utilities		
	OC	LPO	NC	PI	OC	LPO	PI
100	-0.64 (0.37)	-0.85 (0.33)	-0.90 [1] (0.82)	-2.14 [4] (4.05)	-1.11 (0.75)	-1.42 (0.66)	-1.58 (0.49)
200	-0.70 (0.55)	-0.90 (0.37)	-80.11 [3] (756)	-4.68 [62] (9.97)	-1.19 (0.78)	-1.56 (0.71)	-1.90 [18] (0.76)
400	-0.73 (0.67)	-0.91 (0.34)	-5.79 [6] (29.23)	0 [99] (N.A.)	-1.25 (1.00)	-1.60 (0.64)	N.A. [100] (N.A.)
600	-0.59 (0.42)	-0.91 (0.30)	XXXXXXXXXX	0 [99] (N.A.)	-1.11 (0.73)	-1.61 (0.61)	N.A. [100] (N.A.)
p	Model 3						
	Power utility				Mixture of power utilities		
	OC	LPO	NC	PI	OC	LPO	PI
100	-0.78 (1.69)	-1.17 (0.70)	-1.18 (1.94)	-74.31 [23] (487)	-1.23 (0.95)	-1.81 (0.58)	-1.95 [9] (0.42)
200	-0.81 (0.57)	-1.02 (0.16)	-4.18 [11] (16.65)	-3.22 [86] (8.31)	-1.24 (1.04)	-1.87 (0.36)	-1.99 [61] (2.75)
400	-0.70 (0.39)	-1.02 (0.13)	-2.78 [8] (5.89)	0 [98] (0)	-1.16 (0.82)	-1.87 (0.35)	N.A. [100] (N.A.)
600	-0.85 (0.47)	-1.00 (0.11)	XXXXXXXXXX	N.A. [100] (N.A.)	-1.39 (0.95)	-1.83 (0.32)	N.A. [100] (N.A.)

It can be seen from Table 1 that the plug-in method is the clear loser in all cases. It generates the smallest averaged utility values and the largest variance for the terminal utility. It also produces negative wealth, bankrupting the investor. The situation gets worse when the p (or p/n) increases. This is consistent with our assertions regarding the Merton problem degeneracy with traditional implementation, when plugging in the sample estimates.

The NC portfolio that corresponds to the mixture of power utilities is not available in the literature, so it does not show up on the right panel of Table 1. Its applicability therefore seems restrictive in relation to the plug-in and LPO approaches. In the left panel, the NC portfolio does improve significantly compared with the plug-in approach but is still not satisfactory. A negative wealth can still be generated for all correlation matrices for $n < p$ and it no longer works for $p > n$. When it does work, the sample objective value is still much less than that of the oracle portfolio, which contains no estimation errors.

Table 2

Mean (standard deviation) of the RC (defined below) values over 100 simulations.

p	Model 1	Model 2	Model 3
100	1.83 (0.57)	7.10 (1.59)	9.81 (2.55)
200	1.80 (0.41)	7.59 (1.53)	7.34 (0.63)
400	1.73 (0.32)	8.15 (1.64)	6.76 (0.85)
600	1.70 (0.40)	8.29 (1.48)	6.32 (0.49)

The LPO approach is the best of the three implementations. Its performance is comparable to the oracle portfolios. The corresponding variance is even smaller than that of the oracle portfolio, as the LPO approach selects fewer risky assets. This agrees with the proof in the previous section, where the LPO portfolio is as optimally diversified as the oracle but uses a subset of assets for a finite p . This eventually lowers the variation of the LPO portfolio, and for the same reason, the LPO portfolio almost certainly ensures the nonnegativity of the portfolio wealth. Regardless of $n > p$ or $n < p$, LPO portfolios always exist and offer stable performance. The objective values of the LPO approach are close to those of the oracle portfolio. As the establishment of convergence is in the sense of probability, and the rate of convergence is not typically fast, the LPO convergence trend is not obvious from this simulation exercise. We stress that the oracle objective values still suffer from simulation errors. Given the relatively large variance of the oracle portfolio, it is difficult to distinguish the oracle and LPO objective values statistically.

In order to numerically validate Theorem 4.2, Table 2 investigates the rate of convergence of our LPO portfolio with power utility. To this end, we first note that d_n in (4.4) is of order $\sqrt{\log p/n}$ and we examine the error bound of the LPO portfolio scaled by $\sqrt{\log p/n}$, which is given by

$$RC := \left(\frac{\widetilde{V}_{pow}}{V_{pow}} - 1 \right) / \sqrt{\frac{\log p}{n}}.$$

Theorem 4.2 shows that RC should converge to a constant in probability. Table 2 reports the mean and standard deviation of RC s over all simulations for each model and p . We can observe that as p increases, the standard deviation and the increment or decrement of the mean generally decrease. This numerical result supports Theorem 4.2.

6. Empirical studies. Our numerical studies provide us with insights into the finite-sample performance of the LPO portfolios under different correlation structures and the normality assumption of the asset returns distribution. However, the empirical data may not fit all of these. It is important to empirically test the LPO approach with the financial data. Our empirical studies are also out-of-sample.

The aim of this section is to compare the aforementioned portfolios, including the LPO portfolio, with the EW portfolios, which invest in all assets equally. It is suggested in [16] that the EW portfolio has a high out-of-sample performance and should serve as a benchmark

for optimal asset allocation in empirical studies. We consider the following portfolios, with further improvements:

1. AdpLPO refers to the adaptive LPO portfolio that applies the plug-in portfolio to the assets selected by the LPO approach.
2. LPO-EW refers to the portfolio that applies the EW portfolio to the assets selected by the LPO approach.
3. NC-EW refers to the portfolio that applies the EW portfolio to the assets selected by the norm-constrained approach.
4. EWE refers to the portfolio that applies the EW portfolio to $[0.1p]$ assets with the most extreme variance-adjusted returns; see [24].
5. JS refers to the estimated portfolio using the James–Stein estimator of the mean.
6. LPO-JS refers to the LPO portfolio using the James–Stein estimator of the mean.

The first two portfolios are also considered in [35]. The third one is included in [24] for the expected log-utility maximization portfolio. The last two portfolios are similar to PI and LPO portfolios but use the James–Stein estimator of the mean proposed by Jorion [29], which takes the form

$$\begin{aligned}\hat{\beta}_{JS} &= \hat{\mu}_{JS} - r_f \mathbf{1}_p, \quad \mu_{JS} = (1 - \hat{\phi})\hat{\mu}_n + \hat{\phi}\hat{\mu}_{min}, \\ \hat{\phi} &= \frac{p+2}{p+2 + n(\hat{\mu}_n - \hat{\mu}_{min})' \hat{\Sigma}_n^{-1} (\hat{\mu}_n - \hat{\mu}_{min})},\end{aligned}$$

where $0 < \hat{\phi} < 1$, $\hat{\mu}_n$ and $\hat{\Sigma}_n$ are given in (3.1)–(3.2), and $\hat{\mu}_{min} = \hat{\mu}'_n \hat{\Sigma}_n^{-1} \mathbf{1}_p$ is the average excess return on the global minimum-variance portfolio. Notice that $\hat{\Sigma}_n^{-1}$ may not be computable when $n < p$.

6.1. Data description and empirical design. Our empirical studies use the daily historical data of S&P500 components, downloaded from Yahoo Finance for the period January 3, 2005, to December 31, 2014. The S&P500 component list is current as of June 30, 2015. For simplicity, we eliminate the assets with incomplete data. The dataset contains $p = 448$ assets. We conduct two empirical studies with one-year ($n = 252$) and five-year ($n = 1260$) estimation windows, respectively. Consequently, we test the aforementioned portfolios between 2006 and 2014 for the one-year estimation window and between 2010 and 2014 for the five-year window.

We use a power utility with the risk aversion coefficient $\gamma = 2$. We assume that the investor trades his initial wealth $x_0 = 1$ on the S&P500 components over two-week intervals ($n.test = 10$). We use the one-month Treasury constant maturity rate at the beginning of each biweekly investment as a proxy of the risk-free interest rate over the coming two weeks. This choice better reflects the running economic situation. At the beginning of each investment, we use the previous M -year ($M = 1$ or 5) historical data to estimate the parameters in EUM portfolios, or design the norm-constrained portfolios for the coming two weeks. Eventually, we have 25 biweekly investments each year. We record the utilities of those 25 resulting terminal wealths and report the mean and standard deviation of the utilities to evaluate the portfolio performance of the year. Similar to the previous section, we also record the frequency of negative portfolio wealth.

A noteworthy feature of the LPO approach is that it selects a subset of the assets to invest. Hence, the cost of rebalancing a LPO portfolio is less than for other portfolios. To illustrate

this, we introduce a measure, called portfolio turnover,

$$TO = \frac{1}{n.test} \sum_{i=1}^{n.test} \sum_{j=1}^p \left(|\hat{u}_{j,t_i} - \hat{u}_{j,t_i^-}| \right),$$

where $\hat{u}_{j,t_i}$ is the money amount invested in asset j at time t_i , $\hat{u}_{j,t_i^-}$ is the money amount invested in asset j before rebalancing at time t_i , and $\{t_1, \dots, t_{n.test}\}$ are the trading days. The portfolio turnover measures the average wealth traded for a particular strategy. Therefore, the portfolio turnover can be used to measure the possible effects associated with transaction costs, and the smaller the TO , the lower the costs. We report the mean and standard deviation of TO s over all biweekly investments for each portfolio. Our studies abstract the contribution of transaction costs, but the TO may give some insight into their effect.

6.2. Empirical results. We report only the results with a five-year rolling estimation window here, giving the results with a one-year rolling estimation window in Appendix B, as the NC and JC type strategies are infeasible for the one-year scenarios due to the singular sample covariance matrices ($p > n$). We are primarily interested in the comparison between NC and LPO portfolios. However, the one-year rolling estimation window has many practical applications.

Table 3 summarizes the empirical results. It is again clear that the plug-in approach is really very poor, as it generates many negative wealths, produces the smallest objective value, and suffers from a high variance. The norm-constrained strategy again improves the portfolio performance, as shown in [18]. However, compared to the EW, the NC portfolio underperforms

Table 3

Mean (standard deviation) of the terminal values of the power utility over the year (25 times) with the five-year rolling estimation window. The frequency of negative wealth is given in brackets.

Year	2010	2011	2012	2013	2014	TO
EW	-0.993 (0.028)	-1.001 (0.053)	-0.993 (0.026)	-0.988 (0.022)	-0.994 (0.021)	0.00902 (0.00168)
LPO	-0.992 (0.026)	-1.008 (0.062)	-0.998 (0.008)	-0.992 (0.018)	-0.986 (0.083)	0.0123 (0.0332)
NC	-1.001 (0.378)	-1.258 (0.985)	-3.359 (10.669)	-1.494 (1.124)	-1.260 (0.915)	29.230 (129.170)
PI	-13.784 (51.944)	-4.523 [3] (9.538)	-22.005 [3] (72.147)	-7.391 [2] (17.615)	-4.166 [4] (5.878)	173.197 (183.871)
AdpLPO	-0.919 (0.113)	-1.044 (0.209)	-0.988 (0.084)	-0.964 (0.061)	-1.042 (0.534)	0.183 (0.401)
LPO-EW	-0.934 (0.089)	-1.000 (0.044)	-0.992 (0.026)	-0.990 (0.008)	-0.991 (0.023)	0.00520 (0.00497)
NC-EW	-0.990 (0.026)	-1.001 (0.045)	-0.994 (0.025)	-0.988 (0.022)	-0.995 (0.021)	0.00902 (0.00293)
EWE	-0.991 (0.029)	-1.001 (0.054)	-0.995 (0.028)	-0.988 (0.021)	-0.993 (0.020)	0.00879 (0.00172)
JS	-12.949 (48.297)	-4.330 [3] (9.012)	-19.604 [3] (63.037)	-7.028 [2] (16.365)	-4.036 [4] (5.646)	170.940 (180.475)
LPO-JS	-0.992 (0.026)	-1.008 (0.061)	-0.998 (0.008)	-0.992 (0.018)	-0.987 (0.082)	0.0122 (0.0330)

in all cases and suffers from a huge TO , implying that in practice the transaction cost becomes a critical concern.

The LPO approach outperforms the plug-in and NC approaches in all cases, in terms of objective value, variance, and the turnover rate. However, it cannot consistently outperform the EW portfolio. It beats it two out of five times, while in the remaining three cases both strategies perform very similarly. The EW strategy has a smaller turnover rate, and the fact that it is free of estimation error is a well-known advantage.

To further improve the LPO and NC portfolios, we apply the EW strategy to assets selected by these two methods and consider the adaptive LPO approaches previously described and the EWE strategy introduced in [24]. It can be seen from Table 3 that the adaptive LPO manages to beat the EW strategy three out of five times, but the turnover rate also increases. The NC-EW is much improved and has almost the same performance as the EW strategy, in terms of the objective value, variance, and turnover rate. The performance of the NC-EW, EWE, and EW approaches is very similar, as the NC portfolio does not filter out too many stocks from the original sample.

The LPO-EW also improves, beating the EW strategy four out of five times. The turnover rate is also much smaller. In fact, it has the best performance, and though LPO and NC approaches will filter out unfavored stocks, in this empirical study our preference is for the selection from the LPO approach.

To examine other improved estimators rather than the MLEs, we consider the estimated and LPO-JS portfolios using the James–Stein estimator of the mean. Although the JS results are improvements over the PI results, their difference is not significant, so is not the difference between LPO-JS and LPO portfolios. However, we remark here that incorporating the LPO approach with other improved estimators is straightforward. The best candidate estimator is yet to be found and it is interesting to study the corresponding statistical and oracle properties of the improved LPO portfolio.

Our empirical reasons, given in the appendix, are also in favor of the LPO and LPO-EW approaches. We empirically show that the LPO portfolio outperforms other existing data-driven portfolios under the cases of $n < p$ and $n > p$. The rebalancing cost of a LPO portfolio is also relatively low.

7. Conclusion and future works. We have shown that the Merton problem, implemented with the plug-in approach, degenerates to negative infinity in probability for unknown drifts and a known covariance matrix. It degenerates to zero risky investment once the covariance matrix is unknown and the drifts are known. To resolve the degeneracy, we propose a constrained ℓ_1 minimization approach and prove its oracle properties for $\frac{1}{n} \log p \rightarrow 0$. We also demonstrate the wider application of the proposed approach with the mixture of power utilities. Empirical studies show that the LPO approach is better than the norm-constrained and plug-in approaches. It can also be seen from the empirical results that the LPO approach is able to select favorable risky assets according to the expected utility maximization objective. This eventually offers a portfolio strategy that combines the advantages of the Merton problem and the EW portfolio.

Although our discussion is based on the normality assumption on stock returns, the extension to a heavy-tailed distribution seems possible (see [35]) but is nontrivial. A promising

future work may consider jump-diffusion models. According to [1], a suitable choice of the joint jump distribution and the Poisson arrival times decomposes the optimal EUM policy under a jump-diffusion model into a Brownian component and its orthogonal component. The Brownian component is exactly the same as the Merton solution considered in this paper. Hence, the LPO approach can be straightforwardly applied to the Brownian component. The orthogonal component is, however, a low-dimensional solution that has less concern in estimation error. Future research can investigate the degeneracy issue and the oracle property of the LPO framework applied to the jump-diffusion models in [1].

Appendix A. Extension to robust portfolio optimization. A popular topic in portfolio theory is robust portfolio optimization, which takes model uncertainty into account; see [21, 34]. One simple formulation is to consider alternative measures equivalent to the original $\mathbb{P}$: $\mathcal{Q} = \{\mathbb{Q} : \mathbb{Q} \sim \mathbb{P}\}$. By Girsanov's theorem, for each $\mathbb{Q} \in \mathcal{Q}$, there is a stochastic process $\varphi^{\mathbb{Q}}(t)$, which can be regarded as the model misspecification factors, so that

$$\left. \frac{d\mathbb{Q}}{d\mathbb{P}} \right|_{\mathcal{F}_t} = \nu(t) = \exp \left(\int_0^t \varphi^{\mathbb{Q}}(s)' dW_s - \frac{1}{2} \int_0^t \varphi^{\mathbb{Q}}(s)' \varphi^{\mathbb{Q}}(s) ds \right).$$

The value function is then modified with a penalty function for choosing an alternative measure:

$$V(t, x) = \sup_{u \in \mathcal{L}_{\mathcal{F}}^2(0, T; \mathbb{R}^p)} \inf_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}^{\mathbb{Q}} \left[U(X(T)) + \frac{1}{\xi} \int_t^T \frac{P(s)}{\phi(s)} ds \middle| X(t) = x \right],$$

where ξ is a measure of ambiguity aversion, $P(t) := \varphi^{\mathbb{Q}}(t)' \varphi^{\mathbb{Q}}(t)/2$ measures the relative entropy between $\mathbb{P}$ and $\mathbb{Q}$, and $\phi(t)$ is a preference parameter related to the ambiguity aversion, which is suggested by [34] to take the form

$$\phi(t, x) = \frac{1}{R(t, x)V_x(t, x)} = -\frac{V_{xx}(t, x)}{V_x^2(t, x)} \geq 0.$$

Then the optimal controls are given by

$$\varphi^{\mathbb{Q}^*}(t) = -\frac{\xi}{\xi + 1} \sigma' \Sigma^{-1} \beta, \quad u^*(t, x) = \frac{R(t, x)}{\xi + 1} \Sigma^{-1} \beta,$$

and $R(t, x)$ is now driven by

$$R_t + \frac{\Theta_p}{2(\xi + 1)} R^2 R_{xx} + rxR_x - rR = 0, \quad R(T, x) = -U'(x)/U''(x).$$

The LPO approach can therefore be applied to the robust optimal portfolios, and its oracle properties can be obtained without much extra effort.

Once we solve for the risk-tolerance function $R(t, x)$ and estimate $\Sigma^{-1}\beta$ with the LPO approach, the estimated robust optimal portfolios are ready for implementation. In addition to log, power, and exponential utilities, the mixture of power utilities can be considered: see [9, 22, 34]. The guidelines for using the LPO approach are clear and simple, so it is widely applicable.

Table 4

Mean (standard deviation) of the terminal values of the power utility over a year (25 investments) with the one-year rolling estimation window.

Year	2006	2007	2008	2009	2010	
EW	-0.99 (0.02)	-1.00 (0.02)	-1.02 (0.05)	-0.98 (0.06)	-0.99 (0.04)	
LPO	-1.01 (0.11)	-0.98 (0.22)	-1.60 (3.03)	-1.00 (0.01)	-0.99 (0.12)	
PI	-137.05 [20] (302.33)	-2.27 [22] (3.51)	-0.83 [23] (1.16)	-3.56 [22] (4.03)	-5.40 [18] (9.07)	
AdpLPO	-1.09 (0.69)	-1.04 [1] (0.91)	-7.72 [1] (31.29)	-0.98 (0.09)	-1.11 (0.53)	
LPO-EW	-0.87 (0.19)	-0.88 (0.18)	-0.97 (0.09)	-1.00 (0.06)	-0.99 (0.03)	
EWE	-0.99 (0.02)	-0.99 (0.02)	-1.02 (0.05)	-0.98 (0.06)	-0.99 (0.04)	
Year	2011	2012	2013	2014		TO
EW	-1.00 (0.04)	-0.99 (0.02)	-0.99 (0.02)	-0.99 (0.02)		0.011 (0.005)
LPO	-1.05 (0.30)	-1.01 (0.08)	-1.03 (0.20)	-1.01 (0.18)		0.18 (0.59)
PI	NA [25] (NA)	-0.35 [23] (0.41)	-1.88 [21] (3.17)	NA [25] (NA)		4.65e4 (4.15e5)
AdpLPO	-1.06 [1] (0.39)	-1.37 (1.00)	-1.20 (0.53)	-1.69 (2.04)		3.48 (14.01)
LPO-EW	-1.01 (0.05)	-0.99 (0.02)	-1.00 (0.02)	-1.00 (0.03)		0.005 (0.006)
EWE	-1.00 (0.04)	-0.99 (0.03)	-0.99 (0.02)	-0.99 (0.02)		0.011 (0.005)

Appendix B. Supplementary empirical study. Table 4 shows the empirical results using the one-year rolling windows for estimation. In this case, the norm-constrained portfolios are unavailable.

The results for the plug-in portfolios can be ignored because they always lead to bankruptcy. It can be seen that the LPO and EWE portfolios are comparable to the EW portfolio. The LPO-EW portfolio performs slightly better than the EW portfolio, demonstrating that the assets selected by the LPO approach are favorable. The AdpLPO portfolio also improves compared to the plug-in portfolio, but it still has a relatively large variance of the utility compared to others.

The EW portfolio is very robust to economic situations. Hence, it is consistent with the conclusion in [35] that the LPO-EW portfolio empirically performs well, as selecting the key assets using the LPO approach and applying the EW portfolio to them can avoid the effects of estimation errors. During the financial crisis of 2007–2008 and the Greek debt crisis of 2011–2012, the data-driven portfolios can be seen to have performed poorly, as there are obviously structural breaks in those periods. The LPO portfolios in 2009 invest less in risky assets because the CV procedure uses the past one-year (2008) data, which suggests no risky investment. Notice that all data-driven portfolios have the same problem, in that

the historical data may not be able to reflect the future movement of the market. A shorter period of training data would give more stationary information but lead to large estimation errors. This is a reasonable trade-off.

As for the portfolio turnovers, it can be seen that the volatile plug-in portfolio has very large TO . The LPO-EW, EWE, and EW portfolios have minimal TO s, while the TO s of the LPO and AdpLPO portfolio are moderate, as they suffer from estimation errors.

Acknowledgments. We are grateful to seminar participants at the 2014 Quantitative Methods in Finance Conference and the workshop series “Broad Perspectives and New Directions in Financial Mathematics” at the Institute for Pure and Applied Mathematics (IPAM) for helpful comments. We also appreciate the constructive comments from several anonymous referees and the editor.

REFERENCES

- [1] Y. AIT-SAHALIA, J. CACHO-DIAZ, AND T. R. HURD, *Portfolio choice with jumps: A closed-form solution*, Ann. Appl. Probab., 19 (2009), pp. 556–584.
- [2] Y. AIT-SAHALIA, J. FAN, AND D. XU, *High-frequency covariance estimates with noisy and asynchronous financial data*, J. Amer. Statist. Assoc., 105 (2010), pp. 1504–1517.
- [3] F. E. BENTH, K. H. KARLSEN, AND K. REIKVAM, *Merton’s portfolio optimization in a Black-Scholes market with non-Gaussian stochastic volatility of Ornstein-Uhlenbeck type*, Math. Finance, 13 (2003), pp. 215–244.
- [4] P. J. BICKEL AND E. LEVINA, *Some theory for Fisher’s linear discriminant function, ‘naive Bayes’, and some alternatives when there are many more variables than observations*, Bernoulli, 10 (2004), pp. 989–1010.
- [5] M. E. BOCK, *Minimax estimators of the mean of a multivariate normal distribution*, Ann. Statist., 3 (1975), pp. 209–218.
- [6] P. BRANDIMARTE, *Numerical Methods in Finance and Economics: A MATLAB-Based Introduction*, Wiley, New York, 2006.
- [7] M. BROADIE, *Computing efficient frontiers using estimated parameters*, Ann. Oper. Res., 45 (1993), pp. 21–58.
- [8] J. BRODIE, I. DAUBECHIES, C. DE MOL, D. GIANNONE, AND I. LORIS, *Sparse and stable Markowitz portfolios*, Proc. Nat. Acad. Sci. USA, 106 (2009), pp. 12267–12272.
- [9] M. K. BRUNNERMEIER AND S. NAGEL, *Do wealth fluctuations generate time-varying risk aversion? Micro-evidence on individuals*, Amer. Econom. Rev., 98 (2008), pp. 713–736.
- [10] T. T. CAI, W. D. LIU, AND X. LUO, *A constrained l_1 minimization approach to sparse precision matrix estimation*, J. Amer. Statist. Assoc., 106 (2011), pp. 594–607.
- [11] T. T. CAI AND W. D. LIU, *A direct estimation approach to sparse linear discriminant analysis*, J. Amer. Statist. Assoc., 106 (2011), pp. 1566–1577.
- [12] T. T. CAI AND J. LV, *Discussion: The Dantzig selector: Statistical estimation when p is much larger than n* , Ann. Statist., 35 (2007), pp. 2365–2369.
- [13] E. CANDÈS AND T. TAO, *The Dantzig selector: Statistical estimation when p is much larger than n* , Ann. Statist., 35 (2007), pp. 2313–2351.
- [14] V. K. CHOPRA AND W. T. ZIEMBA, *The effect of errors in means, variances, and covariances on optimal portfolio choice*, J. Portfolio Management, 19 (1993), pp. 6–11.
- [15] V. DEMIGUEL, L. GARLAPPI, F. J. NOGALES, AND R. UPPAL, *A generalized approach to portfolio optimization: Improving performance by constructing portfolio norms*, Management Sci., 55 (2009), pp. 798–812.
- [16] V. DEMIGUEL, L. GARLAPPI, AND R. UPPAL, *Optimal versus naive diversification: How inefficient is the $1/N$ portfolio strategy?*, Rev. Financial Stud., 22 (2009), pp. 1915–1953.
- [17] B. DIRIS, F. PALM, AND P. SCHOTMAN, *Long-term strategic asset allocation: An out-of-sample evaluation*, Management Sci., 61 (2015), pp. 2185–2202.

- [18] M. S. DUBOIS AND L. A. M. VERAART, *Optimal diversification in the presence of parameter uncertainty for a risk averse investor*, SIAM J. Financial Math., 6 (2015), pp. 201–241.
- [19] N. EL KAROUI, *High-dimensionality effects in the Markowitz problem and other quadratic programs with linear constraints: Risk underestimation*, Ann. Statist., 38 (2010), pp. 3487–3566.
- [20] J. FAN, J. ZHANG, AND K. YU, *Vast portfolio selection with gross-exposure constraints*, J. Amer. Statist. Assoc., 107 (2012), pp. 592–606.
- [21] J.-P. FOUQUE, C. S. PUN, AND H. Y. WONG, *Portfolio optimization with ambiguous correlation and stochastic volatilities*, SIAM J. Control Optim., 54 (2016), pp. 2309–2338.
- [22] J.-P. FOUQUE, R. SIRCAR, AND TH. ZARIPHPOULOU, *Portfolio optimization & stochastic volatility asymptotics*, Math. Finance (2015), doi:10.1111/mafi.12109.
- [23] L. GARLAPPI, R. UPPAL, AND T. WANG, *Portfolio selection with parameter and model uncertainty: A multi-prior approach*, Rev. Financial Stud., 20 (2007), pp. 41–81.
- [24] A. GANDY AND L. A. M. VERAART, *The effect of estimation in high-dimensional portfolios*, Math. Finance, 23 (2013), pp. 531–559.
- [25] N. J. HIGHAM, N. STRABIC, AND V. SEGO, *Resorting definiteness via shrinking, with an application to correlation matrices with a fixed block*, SIAM Rev., 58 (2016), pp. 245–263.
- [26] K. INADA, *On a two-sector model of economic growth: Comments and a generalization*, Rev. Econom. Stud., 30 (1963), pp. 119–127.
- [27] R. JAGANNATHAN AND T. MA, *Risk reduction in large portfolios: Why imposing the wrong constraints helps*, J. Finance, 58 (2003), pp. 1651–1684.
- [28] W. JAMES AND C. STEIN, *Estimation with quadratic loss*, in Proceedings of the 4th Berkeley Symposium on Probability and Statistics, University of California Press, Berkeley, 1961, pp. 361–379.
- [29] P. JORION, *Bayes-Stein estimation for portfolio analysis*, J. Financial Quant. Anal., 21 (1986), pp. 279–292.
- [30] O. LEDOIT AND M. WOLF, *Honey, I shrunk the sample covariance matrix*, J. Portfolio Management, 30 (2004), pp. 110–119.
- [31] S. LEE, F. ZOU, AND F. A. WRIGHT, *Convergence of sample eigenvalues, eigenvectors, and principal component scores for ultra-high dimensional data*, Biometrika, 101 (2014), pp. 484–490.
- [32] R. C. MERTON, *On estimating the expected return on the market: An exploratory investigation*, J. Financial Econom., 8 (1980), pp. 323–361.
- [33] R. C. MERTON, *Continuous-Time Finance*, Wiley-Blackwell, New York, 1992.
- [34] C. S. PUN AND H. Y. WONG, *Robust investment-reinsurance optimization with multiscale stochastic volatility*, Insurance Math. Econom., 62 (2015), pp. 245–256.
- [35] C. S. PUN AND H. Y. WONG, *High-Dimensional Static and Dynamic Portfolio Selection Problems via ℓ_1 Minimization*, Working paper, Chinese University of Hong Kong, <http://ssrn.com/abstract=2607324> (2015).
- [36] C. STEIN, *Inadmissibility of the usual estimator for the mean of a multivariate normal distribution*, in Proceedings of the 3rd Berkeley Symposium on Probability and Statistics, University of California Press, Berkeley, 1955, pp. 197–206.
- [37] R. VANDERBEI, *Linear Programming, Foundations and Extensions*, Springer, New York, 2008.
- [38] T. ZARIPHPOULOU, *A solution approach to valuation with unhedgeable risks*, Finance Stoch., 5 (2001), pp. 61–82.