Optimal Portfolio Choice with Economic Constraints: An Interpretable Machine Learning Approach *

Yang Liu
Hunan University

Guofu Zhou
Washington University in St. Louis

Current version: August, 2024

*We are grateful to Yi Ding (Discussant) and seminar participants at Hunan University, Tongji University, and Xi'an Jiaotong-liverpool University for very helpful comments. We would also like to thank the conference participants at the 2024 Asian Finance Association Annual Conference and the 2024 Symposium of Big Data, AI, and Econometrics at Xiamen University.

Send correspondence to Guofu Zhou, Olin School of Business, Washington University in St. Louis, St. Louis, MO 63130; e-mail: zhou@wustl.edu; phone: 314-935-6384.

Optimal Portfolio Choice with Economic Constraints:
An Interpretable Machine Learning Approach

Abstract

We develop a new approach to construct the mean-variance efficient portfolio by directly targeting the optimal weight with economic-motivated regularization that incorporates economic constraints to guard against overfitting and enhance interpretability. Instead of using noisy estimators of expected return and covariance matrix, we interpret a portfolio rule as a mapping from historical data to optimal weights and take advantage of the vigorous searching capability of genetic programming (GP) to estimate this weighting function directly. We also propose a trading-frictions-based regularization to control model complexity while preserving interpretability. The out-of-sample Sharpe ratio of our GP approach more than doubles those of existing methods. Beyond portfolio choice, we also leverage portfolio theory to derive a model-implied expected return measure from the GP-optimal weight and find that it subsumes the predictability of other machine learning methods in the cross-section of stock returns. Our study highlights the importance of combining machine learning and economic rationale for interpretable machine learning applications in asset pricing.

JEL Classification: G12, G14, G15

Keywords: Interpretable Machine Learning, Portfolio Optimization, Genetic Programming, Economic Constraints

1. Introduction

Since the seminal paper of Markowitz (1952), the mean-variance (MV) framework has been the major model used in practice in asset allocation and active portfolio management. Theoretically, the mean-variance efficient portfolio has a simple explicit expression that depends only on the expected mean and covariance matrix of asset returns. In practice, however, the MV-framework faces several challenges: (1) $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are very difficult to estimate accurately, and the commonly used "plug-in" portfolio based on the sample mean and sample covariance is shown to generate poor out-of-sample (OOS) performance (e.g., Michaud, 1989); (2) it assumes no market friction and a zero transaction cost; (3) the OOS performance worsens as the asset number increases (e.g., Best and Grauer, 1991; Britten-Jones, 1999; Kan and Zhou, 2007). In particular, the estimation of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ requires a balance panel data of return with a large T, and hence it can not deal with the optimization on the whole market. Treating estimated parameters as true parameters, the plug-in method typically results in a substantial deterioration in out-of-sample portfolio performance.

Various strategies have been proposed in the literature to address the above problems. Early works by Brown (1976) and Bawa et al. (1979) show that the plug-in method is generally outperformed by the Bayesian decision rule under a diffuse prior. Jorion (1986, 1991) proposes the use of a Bayes-Stein estimator instead of the sample mean. DeMiguel et al. (2009) propose a new method based on L_1 and L_2 penalties to mitigate estimation risk. More recently, Ao et al. (2019) introduce a high-dimensional sparse regression method to control portfolio risk and use LASSO to estimate optimal portfolio weights. Kan et al. (2022) propose an optimal combining strategy to mitigate estimation risk. However, these studies are generally built on the Markowitz (1952) framework. Although the regularization methods (such as L_1 and L_2 penalties) improve the accuracy of the covariance matrix estimation, these techniques are generally from a statistical perspective and lack intuitive economic interpretation.

In this paper, we make three major contributions. First, instead of using noisy statistical estimators of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, we propose a new economic-motivated and machine-learning-based solution to directly target the optimal weight with economic constraints, *without* estimating $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$. The intuition is that we interpret a portfolio rule as a *weighting* function mapping historical stock characteristics to the optimal weights. This enables us to solve the optimization problems in the

universe of all stocks. Our model extends the method of Brandt, Santa-Clara, and Valkanov (2009), because their model uses a parametric approach and requires a linear function assumption, and hence faces the model misspecification risk and cannot handle a large number of characteristics. Our method, in comparison, employs powerful genetic programming (GP) to estimate the optimal weighting function non-parametrically in a high-dimensional space without assuming a specific functional form. Due to the unique optimization mechanism and the stochastic nature of GP, we also employ it in conjunction with the popular model ensembling method, which helps to search for the global optima and guard against estimation risk.

The empirical results show that our new method outperforms existing state-of-art methods substantially. For example, the annualized out-of-sample (OOS) Sharpe ratio for our GP portfolio in the universe of the largest 100 stocks in the S&P 500 is up to 1.00, while that of the optimal combination portfolio of Kan et al. (2022) is much lower, at around 0.50. It is important to note that our results are quite robust in several ways. First, we independently train the GP model for hundreds of times to get the empirical distribution of the GP portfolios. The results show that even the poorest 5-percentile of the GP portfolio still strongly dominates the optimal combination portfolio of Kan et al. (2022). Second, GP delivers robust and similar performances across various estimations. For example, the 5-, 50-, and 95-percentile of the GP OOS Sharpe ratios are 0.8789, 0.9754, and 1.1147, respectively, which is clustered in a relatively stable range. Third, our results are robust across different stock universes. Specifically, we compare GP with a large number of the state-of-the-art models in various stock universes, and the results show that GP persistently outperforms all other models we examined. In addition to the largest 100 stocks used by Kan et al. (2022), we also follow Ao et al. (2019) to analyze the universe of the S&P 500 constituents. In particular, Ao et al. (2019) propose a machine-learning approach, maximum-Sharpe-ratio estimated and sparse regression (MAXSER), to construct optimal mean-variance portfolios, and they include almost 20 competing models as benchmarks. We utilize all these models, including MAXSER, as benchmarks for our GP model. The results show that MAXSER indeed strongly outperforms other models, but our GP model produces a much higher monthly Sharpe ratio of 1.089, which almost doubles that of 0.591 for the MAXSER.

Our second contribution is to propose a new economic-motivated regularization method to control model complexity while preserving economic interpretability. We doing so by combining the

flexibility of machine learning with the relevance of economic restrictions. Our economic-motivated regularization method allows for incorporating various economic constraints, such as transaction cost, as penalties into the objective explicitly. For example, if the turnover for a portfolio M, with a cost of C per one USD trading, is $Turn_{M,c}$, then a favourable optimizer M is expected to generate portfolios with $Turn_{M,\lambda_1 C} < Turn_{M,\lambda_2 C}$ for $\lambda_1 > \lambda_2 > 0$. In other words, in response to increasing transaction costs, an intelligent portfolio optimizer learns to reduce the model complexity and decrease the portfolio turnover. However, the transaction cost cannot easily be incorporated into the optimization process in the conventional methods. On the contrary, due to the flexibility of GP, the transaction cost can be easily imposed as a penalty and plays a pivotal role in the optimization of our GP framework.

While several recent studies also consider transaction cost in portfolio evaluation (see, e.g., Novy-Marx and Velikov, 2016; DeMiguel et al., 2020), our transaction cost mitigation method is fundamentally different from these works. Specifically, existing studies are generally based on linear models and simply use the true cost when estimating the model. In comparison, we employ the machine learning method and propose a trading-friction-based regularization method to guard against overfitting while preserving interpretability. In other words, existing studies simply use C to examine how the cost deteriorates the strategy performance. In our framework, however, the strategy itself is changed to minimize the net trading cost. That is, we incorporate $\lambda \times C$ as a regularization into the objective, where the λ is the economic penalty coefficient, and is a hyper-parameter determined by the cross-validation process.

The proposed transaction cost-based penalty λ has intuitive economic and statistical interpretations. In a friction-free world, the model should utilize the new information greedily, resulting in very high portfolio turnovers. However, in the real world, aggressive information utilization is no longer optimal due to the cost associated with trading friction. As such, the penalty coefficient λ balances between the benefit of using new information and the cost implied by economic restrictions. For small λ, GP takes advantage of new information aggressively, but at the cost of inclining to overfitting. As λ increases, GP adaptively learns to behave conservatively, guarding against overfitting. Obviously, too conservative information utilization implied by large λ also deteriorates the OOS performance, as it exposes the model to underfitting and converges to the buy-and-hold strategy for an extremely large value of λ. From a statistical learning perspective,

the optimal λ essentially reflects the well-known bias-variance trade-off. Aggressive information utilization under small λ allows us to capture the new information more efficiently at the cost of producing investment position with higher variance (high turnover), while conservative information utilization under large λ produces more stable portfolio weight, but sacrifices the profitability of utilizing new information.

The third main contribution of this paper is to shed light on the link between portfolio optimization and the cross-section of stock returns. Interestingly, although we focus on the optimal weights, our GP weighting function also has important implications for the cross-section of stock pricing, allowing us to derive a GP-implied expected return measure. The intuition is that, given the Markowitz (1952) optimal weight formula which explicitly depends only on the expected mean and covariance matrix, we can derive a GP-implied expected return measure as a function of our GP weight and an estimator of the covariance matrix. Empirically, we use the historical covariance matrix based on the daily return in the past months to derive the expected return measure. Surprisingly, even with such a noisy covariance matrix estimator, the resultant GP-implied expected return measure still has strong predictive information for future stock returns. For example, the value-weighted decile spread portfolio sorted by this return measure earns a sizable return of 1.34% (t-stat: 4.55). More importantly, controlling for various machine-learning-based expected return measures in Liu et al. (2023), such as LASSO and neural networks, the GP-implied expected return still has significant predictability, but not vice versa. In other words, the cross-section prediction implied by our model outperforms all the major machine learning methods used by Gu, Kelly, and Xiu (2020).

Our paper adds to the relatively small literature on the applications of GP into finance. Neely, Weller and Dittmar (1997) seems the first to apply the GP to find profitable technical trading rules. Allen and Karjalainen (1999) apply the GP to find profitable trading rules to beat the S&P 500 index, though unsuccessfully. Recently, Brogaard and Zareei (2023), with modified algorithms, are able to identify stronger time-series predictability of the S&P 500 index. Ready (2002) also use GP to investigate the profitability of the technical trading rules on DJIA index. In addition, Dempster and Jones (2001) and Dunis, Laws, Middleton, and Karathanasopoulos (2015) apply it to currency and commodities. All these existing studies are about using the GP for time series

prediction.[1] In contrast, our paper is about the optimal portfolio choice, and it also has important implication for the cross-section of stock returns. However, there is a hurdle of applying the GP. It is computational time, which is especially critical in our cross section context which deals with a large number of stocks. Indeed, even on a server with an Intel Xeon E7-8890 and 512 GB memory, the computation time takes weeks for our study. Nevertheless, with increasing computing power each year, the application of the GP in finance will surely increase drastically over time, simply due to the flexibility of the algorithm that it can be used to maximize almost any economic objective.[2]

Our paper also adds to the literature of growing applications of machine learning to finance. Chinco, Clark-Joseph, and Ye (2019), Feng, Giglio, and Xiu (2020), Freyberger, Neuhierl, and Weber (2020), Kozak, Nagel, and Santosh (2020), Bryzgalova, Huang, and Julliard (2022), and Avramov, Cheng, Metzker, and Voigt (2022) provide various dimension reduction analysis, while DeMiguel, Martin-Utrera, Nogales, and Uppal (2020), Gu, Kelly, and Xiu (2020), Chen and Velikov (2020), Chordia, Goyal, and Saretto (2020), Patton and Weller (2020), Avramov, Cheng, and Metzker (2021), Bryzgalova, Pelger, and Zhu (2021), Cong, Tang, Wang, and Zhang (2021), and Cong, Feng, He, and He (2022) focus on portfolio performance. Our paper is related more to the latter by using maximizing the Sharpe ratio, providing the pricing kernel that prices all assets. Our paper is the first to use the powerful GP in the optimal portfolio choice with an economic-motivated penalty, which appears to provide both economic insights and improved performance in general.

The rest of the paper is organized as follows. Section 2 discusses the data and the methodology of our GP model. The next three sections compare GP with various state-of-the-art models, in the the stock universe used by the original paper. Section 3 compares GP with the models in Kan et al. (2022) in the largest 100 stocks. Section 4 conducts the comparison with the models in Ao et al. (2019) in the S&P 500 index constituents. Section 5 compares GP with the methods in Liu et al. (2023) in all stocks. Section 6 concludes.

[1] An exception is Liu et al. (2023) who apply GP in the cross-section of stock returns. We compare this method in Section 3.3..

[2] Nordhaus (2001) shows that the computing power has increased by around 80% per year since 1980.

2. Methodology and data

In this section, we first discuss the baseline GP model for maximizing the Sharpe ratio in portfolio choice. We then incorporate the economic constraints and propose an economic-motivated penalty coefficient to control model complexity and enhance interpretability. Next, we introduce the optimization details, including the search space, hyper-parameter tuning, and ensembling. Last, we discuss the data.

2.1. Baseline model: GP with mean-variance efficiency objective

Consider the standard portfolio choice problem of achieving mean-variance efficiency. Denote the *excess* returns for the N risky assets by $\boldsymbol{r} = (r_1, r_2, ..., r_N)'$. Let $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ be the expected mean and covariance matrix of $\boldsymbol{r}$. The mean-variance efficient portfolio of Markowitz (1952) has an explicit expression, $w^* = \frac{1}{\gamma}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}$, where γ is the scale parameter. As discussed earlier, estimation errors in μ and Σ often make the above portfolio perform poorly in practice.

Here, we use GP to solve the above mean-variance (MV) efficiency problem. Hence, we denote our model as GP_{MV}. One favorable feature of GP_{MV} is that we solve the portfolio optimization problem *without* estimating $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$. The intuition for our method is that we interpret a portfolio rule as a weighting function of historical data (not limited to return data) and take advantage of the powerful searching ability of machine learning to estimate the function non-parametrically in a general function space. Mathematically, denote the historical data at the end of time t as Φ_t. Φ_t can include the historical return data, other characteristics for individual stocks, and macro- or market-level indicators. Our goal is to estimate a weighting function $G(\cdot)$, mapping from historical data of individual risky assets ($\Phi_{i,t}$) to (unscaled) weights on the risky assets. Denote

$$w_{i,t+1} = \frac{G(\Phi_{i,t})}{\sum_i |G(\Phi_{i,t})|}, \tag{1}$$

where $w_{i,t+1}$ is the weight on risky asset i at the beginning of period t+1. Note that $\sum_i |w_{i,t+1}| = 1$, and this indicates that we require a 100% margin on long and short positions and naturally avoid extreme values. Since $w_{i,t+1}$ is essentially determined by $G(\cdot)$, we define the scaling function $\dot{G}(\cdot)$ and re-write Eq. (1) as follows:

$$w_{i,t+1} = \dot{G}(G(\Phi_t)) \equiv \frac{G(\Phi_{i,t})}{\sum_i |G(\Phi_{i,t})|}. \tag{2}$$

For a given $G(\cdot)$, denote the portfolio return in period $t+1$ as R_{t+1}:

$$R_{t+1} = \sum_i w_{i,t+1} r_{i,t+1} = \sum_i \dot{G}(G(\Phi_t)) r_{i,t+1}. \tag{3}$$

Denote $\tilde{r}_{t+1} = \{r_{1,t+1}, r_{2,t+1}, ..., r_{N,t+1}\}$, and then we can define function $F(\cdot)$ and re-write Eq. (3) as

$$F(G(\cdot), \Phi_t, \tilde{r}_{t+1}) \equiv \sum_i \dot{G}(G(\Phi_t)) r_{i,t+1} = R_{t+1}. \tag{4}$$

For GP_{MV}, the objective is to achieve the mean-variance efficiency, namely to maximize the Sharpe ratio (SR) of the portfolio return $\{R_t\}$. Hence, the optimization problem of GP_{MV} can be represented as:

$$\max_{G(\cdot)\in\mathcal{M}} SR\{F(G(\cdot), \Phi_t, \tilde{r}_{t+1})|_{t=0}^{t=T}\}, \tag{5}$$

where $\mathcal{M}$ is the search function space, and $G(\cdot)$ is a function mapping from historical data to (unscaled) weights on the risky assets.

2.2. Incorporate economic constraints

Here, we incorporate the economic constraints of transaction cost and short-sale constraints into the baseline model.

2.2.1. Short-sale constraint

To model short-sale constraint, paralleling to Eq. (2), define

$$w_{i,t+1} = |\dot{G}(G(\Phi_t))| \equiv |\frac{G(\Phi_{i,t})}{\sum_i |G(\Phi_{i,t})|}|. \tag{6}$$

Paralleling to Eq. (4), we can define $\dot{F}(\cdot)$ and denote the portfolio return in period $t+1$ as

$$\dot{F}(G(\cdot), \Phi_t, \tilde{r}_{t+1}) \equiv \sum_i |\dot{G}(G(\Phi_t))| r_{i,t+1} = R_{t+1}. \tag{7}$$

Hence, the portfolio optimization problem with short-sale constraint can be represented as:

$$\max_{G(\cdot)\in\mathcal{M}} SR\{\dot{F}(G(\cdot), \Phi_t, \tilde{r}_{t+1})|_{t=0}^{t=T}\}. \tag{8}$$

2.2.2. Transaction cost

Following Robert, Robert, and Jeffrey (2012), we define the portfolio return net of transaction costs as follows:

$$R_{net,t+1} = \{1 - \sum_i c_{i,t+1}|w_{i,t+2} - w_{i,(t+1)+}|\}(1 + R_{t+1}) - 1, \tag{9}$$

where $c_{i,t+1}$ is a cost level that measures the transaction cost per dollar traded for trading asset i,[3] $w_{i,t+2}$ is the weight on asset i at the beginning of period t+2, and $w_{i,(t+1)+}$ is the weight on asset i at the end of period t+1, and R_{t+1} is the portfolio return without transaction cost. The portfolio turnover in the period t+1 is accordingly defined as:

$$Turn_{t+1} = \sum_i |w_i(t+1) - w_i(t+)|. \tag{10}$$

Given Eq. (2) and Eq. (4), We can re-write Eq. (9) as

$$R_{net,t+1} = \{1 - \sum_i c_{i,t+1}|\dot{G}(G(\Phi_{i,t+1})) - w_{i,(t+1)+}|\}(1 + F(G(\cdot), \Phi_t, \tilde{r}_{t+1})) - 1. \tag{11}$$

Denote $\tilde{c}_{t+1} = \{c_{i,t+1}, ..., c_{N,t+1}\}$. Notice that $w_{i,(t+1)+}$ is determined by $w_{i,t+1}$, namely $\dot{G}(G(\Phi_{i,t}))$, and $\tilde{r}_{t+1}$. Hence, paralleling to Eq. (4), we can define $\ddot{F}(\cdot)$ and re-write Eq. (11) as

$$\ddot{F}(G(\cdot), \Phi_t, \Phi_{t+1}, \tilde{r}_{t+1}, \tilde{c}_{t+1}) \equiv R_{net,t+1}. \tag{12}$$

Then, the portfolio optimization problem with transaction cost $\tilde{c}_{t+1}$ can be represented as:

$$\max_{G(\cdot)\in\mathcal{M}} SR\{\ddot{F}(G(\cdot), \Phi_t, \Phi_{t+1}, \tilde{r}_{t+1}, \tilde{c}_{t+1})|_{t=0}^{t=T}\}. \tag{13}$$

Note that the optimization problem in (5) is a special case of (13) with a zero transaction cost ($c_{i,t} \equiv 0$).

2.3. Economic-motivated penalty

While conventional regularization methods, such as the L_1 and L_2 penalty, are commonly used in recent machine learning asset pricing studies (see, e.g., Freyberger, Neuhierl, and Weber, 2020;

[3]It is easy to apply higher transaction cost for short-sale trading.

Gu, Kelly, and Xiu, 2020), these penalties are not feasible for our GP model due to its non-parametricity. Instead, marrying the flexibility of machine learning with the relevance of economic restrictions, we propose an economic-motivated regularization method to control model complexity and enhance interpretability.

Our intuition is to multiply the true unit transaction cost by a penalty coefficient, denoted as λ, which is a *hyper-parameter* to be selected by cross-validation, to *train* the model. A simple example illustrates the underlying intuition. For any portfolio optimizer M, suppose the turnover for the optimal portfolio under the unite transaction cost of one USD trading, denoted as C, is $Turn_{M,C}$.[4] Suppose $\lambda_1 > \lambda_2 > 0$, then a favourable optimizer M is expected to generate portfolios with $Turn_{M,\lambda_1 C} < Turn_{M,\lambda_2 C}$. In other words, in response to increasing transaction costs (higher penalty coefficient λ), an intelligent portfolio optimizer should learn to reduce the model complexity and decrease the portfolio turnover. However, the transaction cost does not come into the optimization process in the conventional methods. On the contrary, the transaction cost can be easily incorporated and plays a role in the objective in our GP framework. Intuitively, the higher the penalty coefficient λ, the lower the portfolio turnover, the simpler the model, and the more conversation of the model fitting. Hence, the cross-validation process may choose an optimal cost coefficient λ that is a little bit higher than 1 in the *training* sample, to guard against in-sample overfitting.

We now formalize the above intuition mathematically. Adding the penalty coefficient λ to Eq (9), the portfolio return net of transaction costs with coefficient λ is

$$R_{net,t+1} = \{1 - \lambda \sum_i c_{i,t+1} | w_{i,t+2} - w_{i,(t+1)+} | \}(1 + R_{t+1}) - 1. \tag{14}$$

Accordingly, the optimization problem with transaction cost-based regularization is:

$$\max_{G(\cdot)\in\mathcal{M}} SR\{\ddot{F}(G(\cdot), \Phi_t, \Phi_{t+1}, \tilde{r}_{t+1}, \tilde{c}_{t+1}, \lambda)|_{t=0}^{t=T}\}, \tag{15}$$

Compared with (13), (15) has an additional hyper-parameter, the penalty coefficient λ, to govern model complexity.

[4]For now, we assume the unit trading cost is equal for all stocks for simplicity. It is easy to consider the heterogeneity of cost among stocks with different features.

Five points are worth noting for the proposed penalty coefficient λ. Firstly, in terms of its economic interpretation, λ straightforwardly controls the portfolio turnover and hence governs the model complexity. The greater the λ, the more conservative the model fitting. Hence, the reasonable range for λ, by its economic rationale, should be more than 1 to mitigate in-sample overfitting. λ below 1 will deteriorate the OOS performance, as it induces in-sample overfitting. However, a too-large value of λ can also deteriorate the model performance, as it may lead to underfitting and converge to the well-known buy-and-hold strategy for extremely large λ. Hence, there should be an optimal λ, and we employ the standard cross-validation process to determine it. Secondly, since an overly complicated model, implied by pretty high turnovers, is unfavorable in terms of interpretability and OOS performance in the world with trading friction, our model manages to govern the model complex efficiently from an economic perspective. In this sense, the proposed penalty shares the same spirit as the popular regularization techniques of introducing prior knowledge to control complexity for non-parametric models. Indeed, our strategy performance-based penalty is economic-motivated and directly related to portfolio performance and hence easy to interpret economically. Thirdly, there are two special cases of (15). The optimization problem (13) is a special case with the λ=1. In addition, λ with an extremely large value, say 100, corresponds to the well-known buy-and-hold strategy. Fourthly, we only apply λ with values greater than 1 in the *training* and *validation* sample to *train* the GP model guarding against overfitting. In the OOS tests, however, we use the true transaction cost, namely setting λ to 1.

Last but not least, the proposed transaction cost-based penalty can also be interpreted from the perspectives of optimal information utilization and bias-variance trade-off. Consider the information perspectives first. The penalty coefficient λ essentially governs the intensity of the machine utilizing the new information. For small λ, GP takes advantage of new information aggressively, but at the cost of inclining to overfitting. As λ increases, GP adaptively learns to behave conservatively, guarding against overfitting. However, overly conservative information utilization implied by large λ also deteriorates the OOS performance, as it exposes the model to underfitting and converges to the buy-and-hold strategy for an extremely large value of λ. Further, the optimal λ also reflects the famous bias-variance trade-off in the statistic learning literature. Aggressive information utilization under small λ allows us to capture the new information more efficiently at the cost of producing investment position with higher variance (high turnover), while conservative information utilization

under large λ produces more stable portfolio weight, but sacrifices the profitability of utilizing new information.

2.4. Optimization mechanism

In this section, we discuss the optimization mechanism for GP. We start with the representation and search space for solutions and then introduce the optimization process. Next, we discuss the hyper-parameter tuning and ensembling methods to mitigate overfitting.

2.4.1. Representation and search space

In GP, the solution candidates are represented as tree structures and can be encoded as function $G(\cdot)$ mapping from characteristics X to portfolio weights. Each individual $G(\cdot)$ is build of two basic primitives, the *terminal* nodes and *function* nodes. Essentially, the terminal nodes provides the inputs to the GP program, and it includes the input characteristics (X) and some random constants. The *function* nodes comes from a pre-defined function set. Panel A of Figure 1 shows an example of the tree-structure individuals. It consists of two characteristics of X_1 and X_2, a random constant of 1, and two function operators of MULTIPLY ($\times$) and ADD ($+$). It can be coded as a function $G(X) = X_1(X_2 + 1)$. In terms of its economic interpretation, this solution represents such a hypothesis about optimal weight: stocks with greater X_1 should have greater weight. In addition, it also assumes that this effect increases with X_2 by adding an interaction item of X_1 and X_2.

The search space $\mathcal{M}$ is spanned by a large set of functions combining an indicator set and an function set. The indicator set X includes the firm characteristics, which will be discussed soon in section 2.6.. The function set includes both commonly used linear and nonlinear operators, examples of which the linear functions are ADD, MINUS, NEGATIVE, and the nonlinear ones are MULTIPLY, DIVIDE, SIN, COS, ABS, and bool-type operator CMP. This enables GP to exploits both the linear and nonlinear predictability of the characteristics. However, though we do not assume any specific function form for $G(\cdot)$, we limit the maximum of tree depth to 30 for tractability. This still enables a sufficiently large space of millions of candidate solutions, and controls the model complexity and overfitting at the same time.

2.4.2. Optimization process

It is important to examine how the GP selects the individuals to maximize the Sharpe ratio. Different from the common gradient-based method, the optimization of GP is based on the principle of Darwinian natural evolution.

Essentially, GP optimizes the given problem by iteratively producing offspring individuals based on genetic operators and then selecting strong individuals by the natural selection principle. The direction of the evolution is characterized by the fitness function, i.e, the optimization objective, which is the portfolio Sharpe ratio in our case. In particular, after initiating the random individuals in the first generation, GP will calculate their associated Sharpe ratios. Then, to produce new individuals for the next generation, the individuals are randomly selected as parent individuals, with the selection probabilistically biased in favor of the relatively fit individuals with greater Sharpe ratios. Next, the parent individuals are combined by genetic operators, such as crossover and mutation, to create new offspring individuals.

After applying these genetic operators to produce offsprings, GP will evaluate the fitness of these offspring and parent individuals, and those with greater Sharpe ratios will survive as individuals in the next generation. Afterward, successive generations are iteratively generated in the same way, until the generation number exceeds a pre-defined max generation Gen.

2.4.3. Hyper-parameter tuning

There are two important hyper-parameters that control the optimization process of the GP. The first is *Population* (Pop), defined as the number of individuals that GP will generate in each generation. The second is *Generation* (Gen), used to determine the maximum generation that the evolution will iterate. Clearly, the pair $<Pop, Gen>$ characterize the searching depth for GP, and have influence on model performance. Since there is no theoretical criterion for the selection of the pair, we follow the most common approach in the literature and select the hyper-parameters in a validation sample. The validation sample can be interpreted as a simulated OOS sample to learn about model complexity and hence to mitigate overfitting.

To further govern model complexity and enhance interpretability, we also propose an economic-motivated regularization penalty based on transaction cost, and hence there is an additional hyper-

parameter λ in (15). As discussed earlier, λ straightforwardly controls the portfolio turnover and hence governs the model complexity. The greater the λ, the more conservative the model fitting. Hence, λ should be more than 1 to guard against in-sample overfitting. Agian, the optimal λ is determined by the standard cross-validation process.

2.4.4. Ensembling

Following Zhang and Bhattacharyya (2004) and Bhowan, Johnston, Zhang and Yao (2012), we adopt an ensemble approach in training our GP model to improve the model robustness and to mitigate overfitting. In particular, for a given $<Pop, Gen>$, we independently estimate GP five times, and we call this one "RUN". This helps search for the global optima rather than being accidentally trapped by a local optimum. Hence, we get $5 \times Pop$ individuals (or models) in a RUN. Finally, we take the average of the top 5 models with the highest training sample Sharpe ratio as the ensembled model in this RUN.

Because of the stochastic nature of GP, we also conduct the model ensembling at the RUN level. In specific, we estimate the GP models for 100 RUNs, and hence we have 100 GP models. As will be discussed soon, Table 1 reports the empirical performance distribution of these 100 GP models. Table 2 further ensembles the GP model at the RUN level.

2.5. Benchmark methods

In this subsection, we introduce alternative strategies as benchmark methods for our GP model in the empirical tests. These strategies can be used to deal with the estimation risk when forming optimal portfolios, including the invariant optimal portfolio rules proposed by Kan, Wang, and Zhou (2022), the use of the shrinkage covariance matrix estimators of Ledoit and Wolf (2004, 2017), imposing the single factor structure of MacKinlay and Pástor (2000), and imposing no-short-sale constraints. We also discuss ad hoc portfolio rules not derived from the optimization framework. We only briefly introduce the optimal weights and insights of the methods, and please refer to the original papers for technical details.

We define γ as the risk aversion coefficient. w_g is the weight of the global minimum variance (GMV) portfolio, and w_z is the weight of an efficient zero-investment portfolio. N is the number

of risk assets and h is the length of the estimation window for computing μ and Σ, with $h > N$. a follows a standard normal distribution with a mean of 0 and a standard deviation of 1. Besides, ψ is defined as the slope of the asymptote to the ex-ante minimum variance frontier.

2.5.1. Invariant optimal portfolio rules

We first consider the invariant optimal portfolio rules proposed by Kan, Wang, and Zhou (2022). They develop new analytical methods to obtain the exact distribution of the out-of-sample returns and the explicit expression of the expected out-of-sample utility of the plug-in rule.

The first rule is the plug-in rule, that is,

$$\hat{w}_{p,t} = \hat{w}_{g,t} + \frac{1}{\gamma}\hat{w}_{z,t}, \tag{16}$$

with the out-of-sample portfolio return $r_{p,t+1} = \hat{w}'_{p,t} r_{t+1}$.

The second rule is the unbiased rule, obtained by setting $\tilde{c} = (h - N - 1)/h$

$$\hat{w}_{u,t} = \hat{w}_{g,t} + \frac{(h - N - 1)}{\gamma h}\hat{\boldsymbol{w}}_{z,t}. \tag{17}$$

We denote this unbiased portfolio as portfolio u, and its out-of-sample portfolio return is $r_{u,t+1} = \hat{\boldsymbol{w}}'_{u,t}\boldsymbol{r}_{t+1}$.

The third rule is the implementable optimal combining rule, that is,

$$\hat{\boldsymbol{w}}_{\boldsymbol{q},t} = \hat{\boldsymbol{w}}_{g,t} + \frac{g_3\left(\hat{\psi}_t^2\right)}{\gamma}\hat{\boldsymbol{w}}_{z,t}, \tag{18}$$

with the out-of-sample portfolio return $r_{q,t+1} = \hat{w}'_{q,t}\boldsymbol{r}_{t+1}$.

The last portfolio rule is based on a Bayes-Stein estimator developed in Jorion (1986, 1991). We term this portfolio rule as the BS rule:

$$\hat{\boldsymbol{w}}_{BS,t} = \hat{\boldsymbol{w}}_{g,t} + \frac{g_4\left(\hat{\psi}_t^2\right)}{\gamma}\hat{\boldsymbol{w}}_{z,t}, \tag{19}$$

where

$$g_4\left(\hat{\psi}_t^2\right) = \frac{(h - N - 2)^2\hat{\psi}_t^2}{(h+1)(h - N - 2)\hat{\psi}_t^2 + h(N+2)}. \tag{20}$$

2.5.2. Rules with shrinkage covariance matrix estimators

Ledoit and Wolf (2004) introduce a shrinkage estimator which is a linear combination of the sample covariance matrix and the identity matrix,

$$\hat{\Sigma}_t^{LW2004} = (1-\rho_t)\,\hat{\Sigma}_t + \rho_t v_t I_N, \tag{21}$$

where I_N is an $N \times N$ identity matrix, v_t is the shrinkage target which equals to the average of the eigenvalues of $\hat{\Sigma}_t$, and ρ_t is the shrinkage intensity

$$\rho_t = \frac{\text{Min}\left[\frac{1}{h^2}\sum_{s=t-h+1}^{t}\left\|(r_s-\hat{\mu}_t)(r_s-\hat{\mu}_t)' - \hat{\Sigma}_t\right\|^2, \left\|\hat{\Sigma}_t - v_t I_N\right\|^2\right]}{\left\|\hat{\Sigma}_t - v_t I_N\right\|^2} \tag{22}$$

with $\|A\| = \sqrt{\text{tr}\,(AA')/N}$ being the Frobenius norm. When $\hat{\Sigma}_t^{LW2004}$ instead of $\hat{\Sigma}_t$ is plug-in, we obtain a plug-in portfolio with the shrinkage covariance matrix

$$\hat{w}_{p,t}^{LW2004} = \hat{w}_{g,t}^{LW2004} + \frac{1}{\gamma}\hat{w}_{z,t}^{LW2004}, \tag{23}$$

where

$$\hat{\boldsymbol{w}}_{g,t}^{LW2004} = \frac{\left(\hat{\Sigma}_t^{LW2004}\right)^{-1}\mathbf{1}_N}{\mathbf{1}_N'\left(\hat{\Sigma}_t^{LW2004}\right)^{-1}\mathbf{1}_N}, \tag{24}$$

$$\hat{\boldsymbol{w}}_{z,t}^{LW2004} = \left(\hat{\Sigma}_t^{LW2004}\right)^{-1}\left[\hat{\mu}_t - \mathbf{1}_N\left(\hat{\mu}_t'\hat{\boldsymbol{w}}_{g,t}^{LW2004}\right)\right]. \tag{25}$$

In their recent paper, Ledoit and Wolf (2017) propose an improved nonlinear shrinkage estimator of the covariance matrix which is more flexible than the previous shrinkage estimator. The new estimator can be computed as

$$\hat{\Sigma}_t^{LW2017} = \hat{U}_t\hat{D}_t\hat{U}_t', \tag{26}$$

where $\hat{U}_t$ is the orthogonal matrix obtained from a spectral decomposition of $\hat{\Sigma}_t$, and $\hat{D}_t$ is a diagonal matrix $\hat{D}_t = \text{Diag}\left(\hat{d}_1(\lambda_1), \ldots, \hat{d}_N(\lambda_N)\right)$ with $\lambda_1, \ldots, \lambda_N$ being the eigenvalues of $\hat{\Sigma}_t$. For $i = 1, \ldots, N$,

$$\hat{d}_i(\lambda_i) = \frac{1}{\lambda_i\left|s(\lambda_i)\right|^2}. \tag{27}$$

where $s(\lambda_i)$ is an estimator of the Stieltjes (1894) transform of the limiting empirical distribution of sample eigenvalues. The plug-in portfolio with the nonlinear shrinkage covariance matrix can be obtained as

$$\hat{w}_{p,t}^{LW2017} = \hat{w}_{g,t}^{LW2017} + \frac{1}{\gamma}\hat{w}_{z,t}^{LW2017}, \tag{28}$$

where

$$\hat{\boldsymbol{w}}_{g,t}^{LW2017} = \frac{\left(\hat{\Sigma}_t^{LW2017}\right)^{-1}\mathbf{1}_N}{\mathbf{1}_N'\left(\hat{\Sigma}_t^{LW2017}\right)^{-1}\mathbf{1}_N}, \tag{29}$$

$$\hat{\boldsymbol{w}}_{z,t}^{LW2017} = \left(\hat{\Sigma}_t^{LW2017}\right)^{-1}\left[\hat{\mu}_t - \mathbf{1}_N\left(\hat{\mu}_t'\hat{\boldsymbol{w}}_{g,t}^{LW2017}\right)\right]. \tag{30}$$

Explicit expressions of the expected out-of-sample utilities of $\hat{w}_{p,t}^{LW2004}$ or $\hat{w}_{p,t}^{LW2017}$ are not available, and we need to use simulation to evaluate them theoretically.

In addition to the use of the shrinkage covariance matrix estimators, we follow Kan, Wang, and Zhou (2022) to further adjust the exposure to the zero-investment portfolio when constructing the optimal portfolios. Without explicit expressions of the expected out-of-sample utilities, we are unable to derive the exact forms of the optimal combining coefficients when $\hat{\Sigma}_t^{LW2004}$ or $\hat{\Sigma}_t^{LW2017}$ are used. As an alternative, we directly apply the implementable optimal combining coefficient $\hat{c}_t = g_3\left(\hat{\psi}_t^2\right)$. Specifically, we study the following two portfolios:

$$\hat{\boldsymbol{w}}_{q,t}^{LW2004} = \hat{\boldsymbol{w}}_{g,t}^{LW2004} + \frac{g_3\left(\hat{\psi}_t^2\right)}{\gamma}\hat{\boldsymbol{w}}_{z,t}^{LW2004}, \tag{31}$$

$$\hat{\boldsymbol{w}}_{\boldsymbol{q},t}^{LW2017} = \hat{\boldsymbol{w}}_{g,t}^{LW2017} + \frac{g_3\left(\hat{\psi}_t^2\right)}{\gamma}\hat{\boldsymbol{w}}_{z,t}^{LW2017}. \tag{32}$$

2.5.3. Rules with MacKinlay-Pástor single factor structure

MacKinlay and Pástor (2000) adopt a different strategy to deal with the estimation risk. They exploit the implications of an asset pricing model with a single risk factor in the estimation of expected returns. Specifically, they assume the covariance matrix to take the following form:

$$\Sigma = \sigma^2 I_N + a\mu\mu', \tag{33}$$

where a and σ^2 are positive scalars. By imposing such a single factor structure, the estimation risk is reduced because fewer parameters need to be estimated (i.e., instead of μ and Σ, we only need to estimate μ and two scalar parameters a and σ^2).

Under assumption (33), we have

$$\Sigma^{-1} = \frac{1}{\sigma^2}\left(I_N - \frac{a\mu\mu'}{\sigma^2 + a\mu'\mu}\right). \tag{34}$$

Therefore, the weights of the optimal portfolio for the no risk-free asset case are given by

$$w_{MP} = \frac{\left(\sigma^2 + a\mu'\mu\right)\mathbf{1}_N - a\mathbf{1}_N'\mu\mu}{N\left(\sigma^2 + a\mu'\mu\right) - a\left(\mathbf{1}_N'\boldsymbol{\mu}\right)^2} + \frac{1}{\gamma}\frac{N\mu - \left(\mathbf{1}_N'\mu\right)\mathbf{1}_N}{N\left(\sigma^2 + a\mu'\mu\right) - a\left(\mathbf{1}_N'\boldsymbol{\mu}\right)^2} \quad . \tag{35}$$

The first term in the above equation is the GMV portfolio and the second term captures the zero investment portfolio in the setup of MacKinlay and Pástor (2000).

At time t, a, σ^2, and μ are obtained by maximizing the log-likelihood function. The closed-form solution is not available, and we use the semianalytical solution in Tu and Zhou (2011). Plugging the estimates $\hat{a}_{MP,t}, \hat{\sigma}^2_{MP,t}$, and $\hat{\mu}_{MP,t}$ into (35), we obtain the implementable version of the optimal portfolio rule based on the MacKinlay-Pástor (MP) estimators

$$\hat{w}^{MP}_{p,t} = \hat{w}^{MP}_{g,t} + \frac{1}{\gamma}\hat{w}^{MP}_{z,t}, \tag{36}$$

where

$$\hat{\boldsymbol{w}}^{MP}_{g,t} = \frac{\left(\hat{\sigma}^2_{MP,t} + \hat{a}_{MP,t}\hat{\mu}'_{MP,t}\hat{\mu}_{MP,t}\right)\mathbf{1}_N - \hat{a}_{MP,t}\mathbf{1}_N'\hat{\mu}_{MP,t}\hat{\mu}_{MP,t}}{N\left(\hat{\sigma}^2_{MP,t} + \hat{a}_{MP,t}\hat{\mu}'_{MP,t}\hat{\mu}_{MP,t}\right) - \hat{a}_{MP,t}\left(\mathbf{1}_N'\hat{\mu}_{MP,t}\right)^2}, \tag{37}$$

$$\hat{\boldsymbol{w}}^{MP}_{z,t} = \frac{N\hat{\mu}_{MP,t} - \left(\mathbf{1}_N'\hat{\mu}_{MP,t}\right)\mathbf{1}_N}{N\left(\hat{\sigma}^2_{MP,t} + \hat{a}_{MP,t}\hat{\mu}'_{MP,t}\hat{\mu}_{MP,t}\right) - \hat{a}_{MP,t}\left(\mathbf{1}_N'\hat{\mu}_{MP,t}\right)^2}. \tag{38}$$

We term this rule as the MP rule. Note that the MP rule is not consistent unless the assumption of Σ in (33) is true. By imposing such a single factor structure, the estimation risk is reduced, but at the same time, a bias is introduced into the optimal portfolio. Therefore, the MP rule provides a tradeoff between estimation errors and bias.

The expression in (36) suggests that we can also adjust the exposure to the zero-investment portfolio when the factor structure such as (33) is imposed. Therefore, we also examine whether applying $\hat{c}_t = g_3\left(\hat{\psi}^2_t\right)$ can further improve the performance of the MP rule. Specifically, we study the following portfolio

$$\hat{\boldsymbol{w}}^{MP}_{q,t} = \hat{\boldsymbol{w}}^{MP}_{g,t} + \frac{g_3\left(\hat{\psi}^2_t\right)}{\gamma}\hat{\boldsymbol{w}}^{MP}_{z,t}. \tag{39}$$

2.5.4. Rule with no-short-sale constraints

Given the sample mean and the sample covariance matrix, the optimal portfolio with no-short-sale constraints is the solution to the following optimization problem:

$$\begin{array}{ll} \max_w & \boldsymbol{w}'\hat{\mu}_t - \frac{\gamma}{2}w'\hat{\Sigma}_t w \\ \text{s.t.} & \boldsymbol{w}'\mathbf{1}_N = 1, \quad \boldsymbol{w} \geq 0_N \end{array}, \tag{40}$$

where $\mathbf{0}_N$ is an $N \times 1$ vector of zeros. This optimization problem can be readily solved using quadratic programming. We term this portfolio rule with no-short-sale constraints as the NS rule, and denote it as $\hat{\boldsymbol{w}}_{p,t}^{NS}$.

2.5.5. Other rules from portfolio optimization

This section presents three portfolio rules that are also derived from the optimization framework but do not maximize the expected out-of-sample utility.

Sample GMV Portfolio: It is known that the sample mean is an imprecise estimator of the population mean. Some even argue that nothing much is lost in ignoring the mean altogether because the estimation error in the sample mean is so large. As a result, instead of the optimal portfolio, it might be better focusing on the sample GMV portfolio:

$$\hat{\boldsymbol{w}}_{g,t} = \frac{\hat{\Sigma}_t^{-1}\mathbf{1}_N}{\mathbf{1}_N'\hat{\Sigma}_t^{-1}\mathbf{1}_N}. \tag{41}$$

GMV with No-Short-Sale Constraints: Jaganathan and Ma (2003) show that the GMV portfolio, with the no-short-sale constraints in place, performs well. Such portfolio is the solution to the following problem:

$$\begin{aligned} &\min_{w} \boldsymbol{w}'\hat{\Sigma}_t\boldsymbol{w} \\ &\text{s.t. } \boldsymbol{w}'\mathbf{1}_N = 1, \quad \boldsymbol{w} \geq 0_N \end{aligned}. \tag{42}$$

Normalized Kan-Zhou (2007) Three-Fund Rule: In the case in which a risk-free asset is available, Kan and Zhou (2007) derive an optimal three-fund rule under the assumption that the returns of the risky assets are i.i.d. multivariate normally distributed. The portfolio optimally combines the risk-free asset, the sample tangency portfolio, and the sample GMV portfolio. In the case without a risk-free asset, one may think that we can obtain a similar optimal portfolio by normalizing the weights such that the investor is fully invested in risky assets. It can be shown that the normalized optimal three-fund rule takes the following form:

$$\hat{w}_{KZ3,t} = \hat{w}_{g,t} + \left(\frac{1}{\mathbf{1}_N'\hat{\Sigma}_t^{-1}\hat{\mu}_t}\right)\left(\frac{\hat{\psi}_{a,t}^2}{\hat{\psi}_{a,t}^2 + N/h}\right)\hat{w}_{z,t}, \tag{43}$$

which is also a combination of $\hat{\boldsymbol{w}}_{g,t}$ and $\hat{\boldsymbol{w}}_{z,t}$ and belongs to the class as specified in (11). We term this portfolio as the normalized Kan-Zhou (KZ) three-fund rule.

2.5.6. Nonoptimization rules

The $1/N$ Rule: The $1/N$ rule refers to the portfolio strategy with equal weights in the N risky assets, that is, $\boldsymbol{w}_{ew} = \mathbf{1}_N/N$. DeMiguel et al. (2009) compare the $1/N$ rule with various optimal rules derived in the case where a risk-free asset is available and find that the $1/N$ rule outperforms the optimal rules in many empirical data sets. When a risk-free asset is not available, the $1/N$ rule is completely free of estimation errors. Under the i.i.d. assumption, the mean and variance of the out-of-sample portfolio return of the $1/N$ rule are given by

$$\mu_{ew} = \frac{\mathbf{1}_N'\mu}{N}, \quad \sigma_{ew}^2 = \frac{\mathbf{1}_N'\Sigma\mathbf{1}_N}{N^2}. \tag{44}$$

The out-of-sample utility of the $1 = N$ rule is therefore

$$\mathrm{U}_{ew} = \mu_{ew} - \frac{\gamma}{2}\sigma_{ew}^2. \tag{45}$$

Volatility Timing: Although the $1/N$ rule completely ignores the information contained in historical return data, Kirby and Ostdiek (2012) propose two portfolio strategies that use partial sample information and show that they outperform the $1/N$ rule.

The first strategy is the volatility timing strategy. Specifically, the weights of this portfolio are given by

$$\hat{w}_{it} = \frac{\left(1/\hat{\sigma}_{it}^2\right)}{\sum_{j=1}^{N}\left(1/\hat{\sigma}_{jt}^2\right)}, \quad i = 1, 2, \ldots, N, \tag{46}$$

where $\hat{\sigma}_{it}^2$ is the estimated conditional variance on the i-th risky asset at time t. This strategy uses only the information about the conditional variance from the sample and ignores the means and the covariances. We denote this strategy as KO_{VT}.

Kirby and Ostdiek (2012) generalize the above volatility timing strategy to the following:

$$\hat{w}_{it} = \frac{\left(1/\hat{\sigma}_{it}^2\right)^{\eta}}{\sum_{j=1}^{N}\left(1/\hat{\sigma}_{jt}^2\right)^{\eta}}, \quad i = 1, 2, \ldots, N, \tag{47}$$

where $\eta \geq 0$ is a tuning parameter that measures the timing aggressiveness. Kirby and Ostdiek (2012) do not specify how to choose η, but evaluate portfolio performance with different values of η, that is, $\eta = 1, 2$, and 4.

Reward-to-Risk Timing: The second strategy proposed by Kirby and Ostdiek (2012) is the reward-to-risk timing strategy. In addition to the conditional variances, this timing strategy also incorporates the conditional means from the sample. The weights of this portfolio are given by

$$\hat{w}_{it} = \frac{\left(\hat{\mu}_{it}^{+}/\hat{\sigma}_{it}^{2}\right)}{\sum_{j=1}^{N}\left(\hat{\mu}_{jt}^{+}/\hat{\sigma}_{jt}^{2}\right)}, \quad i = 1, 2, \ldots, N, \tag{48}$$

where $\hat{\mu}_{it}^{+} = \text{Max}\left(\hat{\mu}_{i,t}, 0\right)$ with $\hat{\mu}_{it}$ being the maximum likelihood estimator of the conditional mean of the i-th asset at time t. Using only the positive value of the estimated conditional mean avoids short sales. We denote this portfolio as KO_{RT}.

Similarly, Kirby and Ostdiek (2012) define a generalized version of the reward-to-risk timing strategy:

$$\hat{w}_{it} = \frac{\left(\hat{\mu}_{it}^{+}/\hat{\sigma}_{it}^{2}\right)^{\eta}}{\sum_{j=1}^{N}\left(\hat{\mu}_{jt}^{+}/\hat{\sigma}_{jt}^{2}\right)^{\eta}}, \quad i = 1, 2, \ldots, N, \tag{49}$$

where $\eta \geq 0$ is a tuning parameter. Again, they evaluate the portfolio performance for $\eta = 1, 2,$ and 4.

In addition to the maximum likelihood estimator of the conditional mean, Kirby and Ostdiek propose an alternative approach to estimate the conditional mean of risky assets. The second estimator exploits the relation between the first and second moments of returns imposed by linear asset pricing models. For a K-factor model, assuming the factors have identical risk premiums, then the portfolio weights are given by

$$\hat{w}_{it} = \frac{\left(\bar{\beta}_{it}^{+}/\hat{\sigma}_{it}^{2}\right)^{\eta}}{\sum_{j=1}^{N}\left(\bar{\beta}_{jt}^{+}/\hat{\sigma}_{jt}^{2}\right)^{\eta}}, \quad i = 1, 2, \ldots, N, \tag{50}$$

where $\bar{\beta}_{it}^{+} = \text{Max}\left(\bar{\beta}_{it}, 0\right)$ and $\bar{\beta}_{it} = (1/K)\sum_{k=1}^{K}\hat{\beta}_{ikt}$ with $\hat{\beta}_{ikt}$ being the estimated conditional beta of asset i with respect to factor k at time t. We denote the reward-to-risk timing strategy based on $\bar{\beta}_{it}^{+}$ as KO_{BT}.

2.6. Data

Following Liu et al. (2023), the historical characteristics Φ_t consists of 15 variables: the market capitalization (size) and 3 past return-based signals, i.e., R_{-1}, $R_{-12,-2}$, and $R_{-60,-13}$, which correspond to the short-term reversal ($SREV$) of Lehmann (1990), Lo and MacKinlay (1990), momentum (MOM) of Jegadeesh and Titman (1993), and long-term reversal ($LREV$) of DeBondt

and Thaler (1985), respectively. In addition, we also include the 11 price moving average (MA) signals used in Han, Zhou and Zhu (2016), including MAs of lag lengths of 3-, 5-, 10-, 20-, 50-, 100-, 200-, 400-, 600-, 800-, and 1000-days. Following the most recent studies, we normalize each indicator in the cross-section such that it has a mean of zero and a standard deviation of one without loss of generality.

Note that to make the results comparable with existing studies, we examine three stock universes. In Section 3.1., we follow Kan et al. (2022) to contain monthly excess returns of 100 individual stocks with the largest market value from the S&P 500 index constituents. The investment universe is updated annually at the end of February of each year t. In Section 3.2., we follow Ao et al. (2019) to use the stock universe of S&P 500 index constituents. In Section 3.3., we follow Liu et al. (2023) to use all domestic common stocks listed on NYSE, AMEX, and Nasdaq stock markets, and exclude close-end funds, real estate investment trusts, unit trusts, American depository receipts, and foreign stock (or stocks that do not have a CRSP share code of 10 or 11). As the literature typically does, we employ the price filter to exclude the stocks with price below \$5.

In our GP applications below, we split the full sample, from 1945:01 to 2021:12, into three subsamples. The training subsample from 1945:01 to 1980:12 is used to train the machine learning models. The validation subsample from 1981:01 to 1990:12 is used to choose the hyper-parameters in these models.[5] The out-of-sample (OOS) subsample from 1991:01 to 2021:12 is used to evaluate the models' predictive performance.

3. Baseline results

In this section, we compare the baseline GP model, without any economic constraints, with various leading models proposed in three recent studies.[6] We first consider the optimal portfolio combining strategy of Kan et al. (2022), then the linear LASSO-based method of Ao et al. (2019), and last the non-linear machine learning-based model of Liu et al. (2023).

It is important to note that we follow the original papers to use three stock universes to conduct

[5]Following Gu, Kelly, and Xiu (2020), we do not choose cross-validation to maintain the temporal ordering of the data for prediction.

[6]We leave the extended GP model with economic constraints in the next Section 4..

comparisons. These universes differ in stock size and illiquidity, two prominent features of portfolio candidates. Hence, our test has important implications for investment practice and can essentially work as a robustness examination for our GP model.

3.1. Results on the largest 100 stocks in S&P 500 index

We first compare our GP model with the optimal portfolio combining strategy of Kan et al. (2022) in the largest 100 stocks from the S&P 500 index, which represent the largest and the most liquid assets.

3.1.1. Distribution of GP portfolios' Sharpe ratio

As discussed earlier in Section , we estimate the GP model for 100 RUNs under a given $<Pop, Gen>$. Table 1 reports the empirical distribution statistics (mean, median, 1, 5, 95, and 99 percentile) of the annualized Sharpe ratio of these 100 GP portfolios. Panel A reports the training sample results. As the Generation increases from 10 to 40, the averaged Sharpe ratio in the training sample increases monotonically from 1.4089 to 1.7933, which is consistent with the more intense search depth. The averaged Sharpe ratio in the validation sample achieves the highest value of 0.8712 at the Generation of 40. Hence, this parameter and the associated model are selected as the one discussed soon in Table 2 for further OOS tests. However, it is interesting to note that the Sharpe ratios and the returns are stable across different parameters, indicating the smoothness of our model.

3.1.2. Comparison in the largest 100 stocks

Table 2 reports the OOS Sharpe ratio comparison for GP and other models in Kan et al. (2022). The first column reports the annualized Sharpe ratio (SR) in the OOS period. The other columns report the p-values of the one-sided tests on H0: the SR for the model in the corresponding column $>$ the SR for the model in the corresponding row. Note that for GP, the $<Pop, Gen>$ is $< 200, 40 >$.[7] Note again that we estimate the GP model for 100 RUNs. GP_{Median} is the model

[7]This is because, in Table 1, the averaged Sharpe ratio for Generation of 40 in the validation sample is the highest. But the results are robust for other Generation.

with the median of the validation sample SR. GP_{Top5}, GP_{Top10}, GP_{Top15}, GP_{Top20} is the average of the top 5, 10, 15, and 20 models, where the models are sorted by the validation sample SR, respectively.

There are a few points to be made about Table 2. First, the GP models in Panel A earn the highest annualized Sharpe ratio at about 1.0, while the Sharpe ratios of the 4 optimal combining portfolios are much lower at 0.3985, 0.5308, 0.5591 and 0.5263 for $\hat{w}_{q,t}$, $\hat{w}_{q,t}^{LW2004}$, $\hat{w}_{q,t}^{LW2017}$ and $\hat{w}_{q,t}^{MP}$, respectively. The p-values indicate that the SR differences between GP and the combining portfolios are economically and statistically significant. For example, the first row in Panel B shows that the p-values for the SR difference between $\hat{w}_{q,t}$ and the five GP portfolios in Panel A are 0.01, 0.01, 0.01, 0.00, and 0.00, respectively.

Second, the ensembling of GP model helps to mitigate estimation risk and further enhance the model performance. In specific, GP_{Median} is a single GP model, while the other 4 models in Panel A are ensembled GP models. As the number of models for ensembling increases to 20, the Sharpe ratio increases from 1.0345 to 1.1001.[8]. The p-values indicate that such improvement is statistically significant as the p-values in Panel A in the column of GP_{Top20} is 0.05 and 0.07, 0.01 and 0.08, when GP_{Top20} is compared with GP_{Median}, GP_{Top5}, GP_{Top10}, and GP_{Top15}, respectively.

Third, the nonoptimization rules in Panel G outperform the 4 optimal combining portfolios but are dominated by the GP portfolio. On average, the nonoptimization rule generates a Sharpe ratio at about 0.70, which is about 30% lower than that of the GP portfolio. In particular, the p-values in the column of GP_{Top20} in Panel G is 0.06, 0.10, 0.07, and 0.06, indicating that GP_{Top20} strongly outperforms the nonoptimization rules. On the other hand, the optimal combining portfolios earn a lower average Sharpe ratio at about 0.50 and the associated p-values in Panel G indicate no statistically significant difference among the performances for optimal combining rules and nonoptimization rules.

[8]If the number of models for ensembling further increases to 25 or 30, the model performance remains almost constant.

3.2. Results on S&P 500 index constituents

Here, we further compare our GP model with the LASSO-based optimization method of Ao et al. (2019) in the S&P 500 index constituents.[9]

3.2.1. Additional benchmark models

Table 3 lists the benchmark portfolios and their abbreviations. In particular, Ao et al. (2019) propose a new approach named maximum-Sharpe-ratio estimated and sparse regression (MAXSER) to construct optimal mean-variance portfolios. They also consider several models with economic constraints and with state-of-the-art techniques to mitigate overfitting as benchmarks. We utilize all these models as benchmarks for our GP model. It is important to note the notation for the GP models:

- "GP-NoCost" stands for the GP model without short-sale constraint and without transaction cost;
- "GP-WithCost" stands for the GP model without short-sale constraint and with transaction cost;
- "GP-NoCost-SS" stands for the GP model with short-sale constraint and without transaction cost;
- "GP-WithCost-SS" stands for the GP model with short-sale constraint and with transaction cost.

As mentioned earlier, in this section, we only examine the baseline GP model without transaction cost (i.e., GP-NoCost and GP-NoCost-SS). We leave those with transaction cost (i.e., GP-WithCost and GP-WithCost-SS) in Section 4..

3.2.2. Comparison in the S&P 500 constituents

Table 4 reports the OOS performance in the in the S&P 500 constituents. Our GP substantially outperforms other models by doubling their Sharpe ratios. When there is no short-sale constraint, Panel A and Panel B show that our GP portfolio (GP-NoCost) generates the greatest annualized Sharpe ratio up to 1.089, almost doubling that of the second-highest annualized Sharpe ratio of 0.591 produced by MAXSER of Ao et al. (2019) . Incorporating the short-sale constraint, Panel C and Panel D show that our GP portfolio (GP-NoCost-SS) again yields the greatest annualized

[9]We report results on the OOS period from 1997:01 to 2016:12 to compare the results with Ao et al. (2019). The results of MAXSER and other models come from Table 5 and Table 6 of Ao et al. (2019). We will update the results to the recent sample.

Sharpe ratio of 0.637, doubling that of the second-highest annualized Sharpe ratio of 0.316 produced by MV-P-SSCV.

3.3. Results on all stocks

Liu et at. (2023) also employ GP to optimize the Sharpe ratio of the usual value-weighted spread portfolio in the universe of all stocks. Their intuition is to maximize the Sharpe ratio of the spread portfolio sorted by the expected return generated by GP. Put differently, Liu et at. (2023) try to find the optimal *pricing* function mapping the past characteristics to expected returns, while we focus on the optimal *weighting* function mapping past characteristics to portfolio weights in (2) and (5).

In this section, we compare these two GP models. Note that we denote GP_{MV} or GP as the GP model with the mean-variance efficiency objective in this paper, and GP_{Spread} the GP model which maximizes the Sharpe ratio of the high-minus-low spread portfolio of Liu et at. (2023). To make the results comparable with Liu et al. (2023), the stock universe in this section includes all domestic common stocks listed on NYSE, AMEX, and Nasdaq stock markets.

3.3.1. GP performance in the full stock sample

Table 5 reports the in-sample and OOS performance of the GP portfolio constructed by Eq. (2) and Eq. (5) under various parameteres. We report the results for various Generation parameters of 10, 20, and 40, and the Population is set to 200. On average, the annualized Sharpe ratio for the training, validation, and OOS periods is about 3.60, 3.50, and 2.00, respectively. It is important to note that the results are pretty robust across the parameters. For example, the OOS Sharpe ratio is 2.014, 2.030, and 1.995 across the Generations. Panel A shows that since the validation Sharpe ratio achieves the highest value of 3.564 for Generation of 10, this parameter is selected as the one used in the following OOS tests.

3.3.2. Comparison with other machine-learning methods

Table 6 compares the optimal portfolio proposed in this paper with the spread portfolios generated by machine learning methods in Liu et at. (2023). The first 11 columns report the summary

statistics for the decile spread portfolio generated by GP (GP_{Spread}) and other machine learning models, and the results are from Table 1 of Liu et at. (2023). The last column reports the GP portfolio constructed by Eq. (2) and Eq. (5).

GP_{MV} earns similar returns with GP_{Spread} of about 1.50%, but its volatility is much lower at 2.55%. As a result, GP_{MV} produces the highest SR of 2.01, much higher than that of GP_{Spread} (1.21), and more than doubling that of other machine learning methods.

4. Effect of the cost penalty coefficient

While Section 3. focuses on the baseline GP model without transaction cost in various stock universes, here, we extend the model to examine the effect of economic constraints on portfolio performance. In particular, we are interested in the effect of the economic-motivated and transaction-cost-based penalty coefficient λ proposed in Section 2.3.. Does λ enhance the model performance as it is designed to? If so, what is the mechanism? We start with a special case with λ of 1 and then compare the performance under different λ. Note that in this section, the stock universe is the S&P 500 constituents, as it allows sufficient diversification on large and liquid stocks.

4.1. Cost penalty coefficient of 1

Table 7 reports the OOS performance of the GP model in the optimization problem (15), but setting the cost penalty coefficient λ constant to 1.[10] Alternatively, it can be viewed as the optimization problem (13) where we do not include the cost penalty coefficient λ. Here, we assume a constant $\tilde{c}_{t+1}$ of 0.1% in (15).

Our GP portfolios still significantly outperform other models by doubling their Sharpe ratios. When there is no short-sale constraint, Panel A and Panel B show that our GP portfolio (GP-WithCost) generates the greatest annualized Sharpe ratio up to 0.794, much higher than the second-highest annualized Sharpe ratio of 0.483 produced by MAXSER. In terms of the case with short-sale constraint, Panel C and Panel D show that our GP portfolio (GP-WithCost-SS) again yields the greatest annualized Sharpe ratio of 0.594, tripling that of the second-highest annualized Sharpe

[10]The next subsection 4.2. will soon compare the OOS performance of GP models estimated with different λ in (15).

ratio of 0.174 produced by MV-P-SSCV.

4.2. Varying cost penalty coefficient

Here, we examine the effect of transaction cost-based penalty coefficient λ by comparing the OOS performance of the GP models estimated with different λ in (15).

As discussed previously in Section 2.3., λ straightforwardly controls the portfolio turnover and hence governs the model complexity. The lower the λ, the more greedy the information utilization, the more aggressive the model fitting, and hence the greater the portfolio turnover. We also assert that λ below 1 will deteriorate the OOS performance as it encourages overly aggressive information utilization which leads to higher portfolio turnover and model overfitting. To verify this hypothesis, we consider two special cases of λ in (15): 0 and 1.

Note again that for a given λ, we estimate the GP model according to (15). Hence, for λ=1, we estimate the GP model in the training sample assuming the transaction cost is the true cost $(1 \times \tilde{c}_{t+1})$. Similarly, for λ=0, we estimate the GP model in the training sample without transaction cost $(0 \times \tilde{c}_{t+1})$. Assume the "true" unit transaction cost $\tilde{c}_{t+1}$ in (15) is 0.1%. Based on our hypothesis, we expect to see that the GP model estimated with smaller λ (i.e., 0) should have a greater turnover and weak Sharpe ratio in the OOS test.[11]

Table 8 verifies our hypothesis. Panel A of Table 8 reports the results with short-sale constraints. Indeed, GP model estimated with λ of 0 produces a higher monthly turnover of 181.0% while that for the model estimated with λ of 1 generates a lower turnover of 164.8%. Further, the reduction in turnover indeed translates into greater OOS performance. For example, GP model estimated with λ of 1 earns higher return, lower volatility, and hence higher Sharpe ratio. Panel B reveals consistent results when the short-sale is allowed.

The above empirical results confirm that our economic-motivated cost penalty parameter λ can indeed improve the OOS performance. We trace the enhanced performance to the fact that parameter λ successfully controls model complexity and guards against overfitting. Moreover, our results can also be interpreted from the perspective of optimal intensity for information utilization and

[11] Note that both GP models (estimated with λ=0 or 1) will be evaluated in the OOS test based on the "true" unit transaction cost $\tilde{c}_{t+1}$. In other words, the λ is always 1 in the OOS.

the bias-variance trade-off. Aggressive information utilization under small λ allows us to capture the new information more efficiently at the cost of producing investment position with higher variance (high turnover), while conservative information utilization under large λ produces more stable portfolio weight (low turnover), but sacrifices the profitability from utilizing new information. Our method can achieve the optimal information utilization intensity and the bias-variance trade-off by selecting an optimal λ in the standard cross-validation method.

Last, it is important to note that although we verifiy that λ indeed affects the performance in the way we expect, the above results are preliminary. First, since we now focus on the large and liquid stocks, we assume a small transaction cost $\tilde{c}_{t+1}$, and hence the effect of λ seems a little bit weak. We expect that the effect should be stronger among stocks that are more expensive to trade. Second, we only test two special cases of λ in (15): 0 and 1, among which 0 is by default an undesirable candidate. As discussed previously in Section 2.3., the candidate for λ should be greater than 1 to mitigate overfitting. In the coming-soon new version, we will try alternative stock universes and set λ to various values, repeat the above analysis and choose the optimal λ by cross-validation.

5. Link to the cross-section of stock returns

While our GP model (GP_{MV}) targets the optimal *weight* function, we also investigate an interesting question: can it also shed light on the *pricing* function of the cross-section of stock returns? Surprisingly, we show that our GP model indeed has important implications for asset pricing by verifying the strong predictability of a GP-implied expected return measure. Note that the stock universe in this section is the full stock sample, as most studies in cross-section return aim to find pervasive pricing patterns and hence focus on the whole sample.

5.1. GP-implied expected return measure

The Markowtiz (1952) optimal weight formula shows $w^* = \Sigma^{-1}\mu$, where μ and Σ is the expected return and covariance matrix. GP_{MV} directly estimates the portfolio weight in Eq. (1), denoted as $\boldsymbol{w}_{MV}$, without estimating μ and Σ. However, based on the $\boldsymbol{w}_{MV}$, we can derive a GP_{MV} implied

expected return by the Markowtiz (1952) formula as

$$ER_{GP-MV} = \hat{\Sigma}\boldsymbol{w}_{MV}, \tag{51}$$

where $\hat{\Sigma}$ is some estimator for the covariance matrix.

We use three covariance matrices to proxy $\hat{\Sigma}$, that is the historical covariance matrix based on the daily return in the past 3, 6, and 12 months (i.e., $\hat{\Sigma}_{3Mon}$, $\hat{\Sigma}_{6Mon}$, and $\hat{\Sigma}_{12Mon}$).

5.2. Predictability of cross-section returns

Here, we examine the cross-section pricing ability of this GP_{MV} implied expected return ER_{GP-MV}. In particular, the pricing function in GP_{Spread} of Liu et al. (2023) directly estimates the expected return, which we denote here as $ER_{GP-Spread}$. Hence, it is interesting to compare the pricing ability of ER_{GP-MV} with $ER_{GP-Spread}$.

Table 9 reports the results. Panel A shows that the portfolios sorted by ER_{GP-MV} exhibit a monotonically increasing return pattern, and the resultant spread portfolio produces sizable returns. In specific, the spread return is 1.34%, 1.48% and 1.47% per month with large t-stats for $\hat{\Sigma}_{3Mon}$, $\hat{\Sigma}_{6Mon}$, and $\hat{\Sigma}_{12Mon}$, respectively. These values are comparable to the spread return of 1.53% for $ER_{GP-Spread}$ previously shown in Table 6.

Panel B and Panel C further compare the pricing ability of ER_{GP-MV} and $ER_{GP-Spread}$ by regressing one on another one in the cross-section, and then examining the spread portfolio sorted by the resultant residuals. Panel B shows that controlling for $ER_{GP-Spread}$, ER_{GP-MV} still generates significant spread returns of 0.90%, 0.98%, ans 1.02% per month for the three $\hat{\Sigma}$. On the other hand, Panel C reports that controlling for ER_{GP-MV}, the portfolio formed by $ER_{GP-Spread}$ shows a flat return pattern and the resultant spread return shrinks to zero. This results indicate that the predictability of $ER_{GP-Spread}$ is subsumed by ER_{GP-MV}.

Table 10 further compares the GP_{MV}-implied expected return with other expected return measures generated by other machine learning methods, such as LASSO and neural networks. Panel A shows that controlling for other models, the spread portfolio sorted by the GP_{MV}-implied expected return persistently yields positive and highly significant returns. On the contrary, Panel B reports that controlling for ER_{GP-MV}, the spread portfolio returns of other machine learning methods are

generally insignificant.

6. Conclusion

We propose a new approach to construct the optimal mean-variance portfolio *without* estimating the expected return and covariance matrix. The intuition is to interpret a portfolio rule as a mapping from historical data to the optimal weights, and we take advantage of the powerful searching capability of genetic programming (GP) to estimate this weighting function non-parametrically. The out-of-sample Sharpe ratio of our GP approach more than doubles those of existing methods, including the optimal combining strategy and other machine learning-based models, such as LASSO and neural networks. We also derive a model-implied expected return measure from the GP-optimal weight, and find that it subsumes the predictability of all major machine learning methods used by Gu, Kelly, and Xiu (2020), such as neural networks, in the cross-section of stock returns.

We also propose an economic-motivated regularization method by incorporating economic constraints as the penalty into the objective to control the model complexity and enhance interpretability. Empirical results confirm that our economic-motivated cost penalty can indeed improve the OOS performance. We trace the enhanced performance to the fact that penalty coefficient successfully controls model complexity and guards against overfitting. Moreover, our regularization method can also be interpreted from the perspective of optimal intensity for information utilization and the bias-variance trade-off. Aggressive information utilization under small regularization parameter λ allows us to capture the new information more efficiently at the cost of producing investment position with higher variance (high turnover), while conservative information utilization under large λ produces more stable portfolio weight (low turnover), but sacrifices the profitability from utilizing new information.

In short, our paper is the first to use the powerful GP in the optimal portfolio choice with an economic-motivated penalty, which provides both economic insights and improved performance in general. Our results indicate that combining the flexibility of machine learning and the relevance of economic restriction is the key to economically interpretable machine learning applications in asset pricing. Our approach provides what appears the most efficient portfolio to date for a given set of assets, which is the most relevant for investment management and the stochastic discount

factor that prices all assets.

References

Allen, F., Karjalainen, R., 1999. Using genetic algorithms to find technical trading rules. Journal of Financial Economics 51, 245-271.

Ao, M., Yingying, L., Zheng, X., 2019. Approaching mean-variance efficiency for large portfolios. Review of Financial Studies 32, 2890-2919.

Avramov, D., Cheng, S., Metzker, L., Voigt, S., 2022. Integrating factor models. Journal of Finance, forthcoming.

Avramov, D., Cheng, S., Metzker, L., 2021. Machine learning versus economic restrictions: Evidence from stock return predictability. Management Science, forthcoming.

Barillas, F., Shanken, J., 2017. Which alpha?. Review of Financial Studies 30, 1316-1338.

Best, M.J., Grauer, R.R., 1991. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results. Review of Financial Studies 4, 315-342.

Bhowan, U., Johnston, M., Zhang, M., Yao, X., 2012. Evolving diverse ensembles using genetic programming for classification with unbalanced data. IEEE Transactions on Evolutionary Computation, 17, 368-386.

Britten-Jones, M., 1999, The sampling error in estimates of mean-variance efficient portfolio weights, Journal of Finance 54, 655–671.

Brogaard, J., Zareei, A., 2023. Machine learning and the stock market. Journal of Financial and Quantitative Analysis 58, 1431-1472.

Bryzgalova, S., Huang, J., Julliard, C., 2022. Bayesian solutions for the factor zoo: We just ran two quadrillion models. Journal of Finance, forthcoming.

Bryzgalova, S., Pelger, M., Zhu, J., 2021. Forest through the trees: Building cross-sections of stock returns. Available at SSRN 3493458.

Chen, A. Y., Velikov, M., 2020. Zeroing in on the expected returns of anomalies. Journal of Financial and Quantitative Analysis, forthcoming.

Chinco, A., Clark-Joseph, A.D., Ye, M., 2019. Sparse signals in the cross-section of returns. Journal of Finance 74, 449-492.

Chordia, T., Goyal, A., Saretto, A., 2020. Anomalies and false rejections. Review of Financial Studies 33, 2134–2179.

Cong, L. W., Feng, G., He, J., He, X., 2022. Growing the Efficient Frontier on Panel Trees. Working paper.

Cong, L. W., Tang, K., Wang, J., Zhang, Y., 2021. Alphaportfolio: Direct construction through reinforcement learning and interpretable AI. Working paper.

Daniel, K., Hirshleifer, D., Sun, L., 2020. Short-and long-horizon behavioral factors. Review of Financial Studies 33, 1673-1736.

Darwin, C., 2004. On the origin of species, 1859. Routledge.

DeBondt, W.F.M., Thaler, R., 1985. Does the stock market overreact? Journal of Finance 40, 783-805.

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

DeMiguel, V., Martin-Utrera, A., Nogales, F. J., Uppal, R., 2020. A portfolio perspective on the multitude of firm characteristics. Review of Financial Studies 33, 2180–2222.

Dempster, M.A. and Jones, C.M., 2001. A real-time adaptive trading system using genetic programming. Quantitative Finance, 1, 397-413.

Dunis, C.L., Laws, J., Middleton, P.W., Karathanasopoulos, A., 2015. Trading and hedging the corn/ethanol crush spread using time-varying leverage and nonlinear models. The European Journal of Finance, 21, 352-375.

Fama, E.F., French, K.R., 1993. Common risk factors in the returns on stocks and bonds. Journal of financial economics 33, 3-56.

Fama, E.F., French, K.R., 2015. A five-factor asset pricing model. Journal of financial economics 116, 1-22.

Fama, E. F., MacBeth, J. D. 1973. Risk, return, and equilibrium: Empirical tests. Journal of political economy, 81(3), 607-636.

Feng, G., Giglio, S., Xiu, D., 2020. Taming the factor zoo: A test of new factors. Journal of Finance, 75, 1327-1370.

Filippou, I., Rapach, D., Taylor, M.P., Zhou, G., 2020. Exchange Rate Prediction with Machine Learning and a Smart Carry Portfolio. Available at SSRN 3455713.

Freyberger. J., Neuhierl A., Weber, M., 2020. Dissecting characteristics nonparametrically, Review of Financial Studies, 33, 2326-2377.

Green, J., Hand, J.R., Zhang, X.F., 2017. The characteristics that provide independent information about average us monthly stock returns. The Review of Financial Studies, 30, 4389-4436.

Gu, S., Kelly, B., Xiu, D., 2020. Empirical asset pricing via machine learning. Review of Financial Studies, 33(5), 2223-2273.

Guo, X., Lin, H., Wu, C., Zhou, G., 2022. Investor Sentiment and the Cross-Section of Corporate Bond Returns. Available at SSRN 3223846.

Han, Y., He, A., Rapach, D. and Zhou, G., 2021. Expected stock returns and firm characteristics: E-LASSO, assessment, and implications. Available at SSRN 3185335.

Han, Y., Zhou, G., Zhu, Y. 2016. A trend factor: Any economic gains from using information over investment horizons?. Journal of Financial Economics 1222, 352-375.

Hou, K., Xue, C., Zhang, L., 2015. Digesting anomalies: An investment approach. Review of Financial Studies 28, 650-705.

Jagannathan, R., Ma, T., 2003. Risk reduction in large portfolios: Why imposing the wrong constraints helps. The Journal of Finance 58, 1651-1683.

Jegadeesh, N., Titman, S., 1993. Returns to buying winners and selling losers: implications for stock market efficiency. Journal of Finance 48, 65-91.

Jiang, J., Kelly, B. T., Xiu, D, 2020. (Re-)Imag(in)ing price trends. Available at SSRN 3756587.

Jorion, P., 1986. Bayes-Stein estimation for portfolio analysis. Journal of Financial and Quantitative analysis 21, 279-292.

Jorion, P., 1991. Bayesian and CAPM estimators of the means: Implications for portfolio selection. Journal of Banking & Finance 15, 717-727.

Kan, R., Wang, X., Zhou, G., 2022. Optimal portfolio choice with estimation risk: No risk-free asset case. Management Science 68, 2047-2068.

Kan, R., Zhou, G., 2007. Optimal portfolio choice with parameter uncertainty. Journal of Financial and Quantitative Analysis 42, 621-656.

Kan, R., Zhou, G., 2012. Tests of mean-variance spanning. Annals of Economics and Finance 13, 139-187.

Kirby, C., Ostdiek, B., 2012. It's all in the timing: simple active portfolio strategies that outperform naive diversification. Journal of Financial and Quantitative Analysis 47, 437-467.

Koza, J.R., 1992. Genetic Programming: On the Programming of Computers by Means of Natural Selection. MIT Press, Cambridge.

Kozak, S., Nagel, S., Santosh, S., 2020. Shrinking the cross-section. Journal of Financial Economics, 135, 271-292.

LeBaron, B., Arthur, W.B., Palmer, R., 1999. Time series properties of an artificial stock market. Journal of Economic Dynamics and Control, 23, 1487-1516.

Ledoit, O., Wolf, M., 2004. A well-conditioned estimator for large-dimensional covariance matrices. Journal of Multivariate Analysis 88, 365-411.

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

Ledoit, O., Wolf, M., 2017. Nonlinear shrinkage of the covariance matrix for portfolio selection: Markowitz meets Goldilocks. The Review of Financial Studies 30, 4349-4388.

Lehmann, B.N., 1990. Fads, martingales and market efficiency. Quarterly Journal of Economics 105, 1-28.

Lewellen, J. 2015. The Cross-section of Expected Stock Returns. Critical Finance Review 4, 1-44.

Liu, Y., Zhou, G., Zhu, Y., 2023. Maximizing the Sharpe Ratio: A Genetic Programming Approach. Available at SSRN 3726609.

Lo, A.W., MacKinlay, A.C., 1990. When are contrarian profits due to stock market overreaction? Review of Financial Studies 3, 175-205.

MacKinlay, A.C., Pástor, Ľ., 2000 Asset pricing models: Implications for expected returns and portfolio selection. The Review of financial studies 13, 883-916.

Markowitz, H., 1952, Portfolio selection, Journal of Finance 7, 77–91.

Michaud, R.O., 1989. The Markowitz optimization enigma: Is 'optimized'optimal?. Financial analysts Journal 45, 31-42.

Neely, C., Weller, P., Dittmar, R., 1997. Is technical analysis in the foreign exchange market profitable? A genetic programming approach. Journal of financial and Quantitative Analysis, 32, 405-426.

Newey, W.K., West, K. D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703-708.

Nordhaus, W.D., 2001. The progress of computing. Available at SSRN 285167.

Novy-Marx, R. Velikov, M., 2016. A taxonomy of anomalies and their trading costs. Review of Financial Studies 29, 104-147.

Patton, A. J., Weller, B., 2020. What you see is not what you get: The costs of trading market anomalies. Journal of Financial Economics 137, 515–549.

Polachek, S.W., Das, T., Thamma-Apiroam, R., 2015. Micro-and macroeconomic implications of heterogeneity in the production of human capital. Journal of Political Economy, 123, 1410-1455.

Ready, M.J., 2002. Profits from technical trading rules. Financial Management, 43-61.

Robert, E., Robert, F., Jeffrey, R., 2012. Measuring and modeling execution cost and risk. Journal of Portfolio Management 38, 14-28.

Schwert, G.W., 2003. Anomalies and market efficiency. In: Constantinides, G.M., Harris, M., Stulz, R.M. (Eds.), Handbook of the Economics of Finance, 1. Elsevier, Amsterdam, Netherlands, pp. 939-974. chap. 15.

Stambaugh, R.F., Yuan, Y., 2016. Mispricing factors. The Review of Financial Studies 30, 1270-1315.

Stieltjes, T.J., 1894. Recherches sur les fractions continue. Annales de la Faculté des sciences de Toulouse: Mathématiques 8, J1-J122.

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

Winschel, V., Krätzig, M., 2010. Solving, estimating, and selecting nonlinear dynamic models without the curse of dimensionality. Econometrica, 78, 803-821.

Zhang, X.F., 2006. Information uncertainty and stock returns. Journal of Finance, 61, 105-137.

Zhang, Y., Bhattacharyya, S., 2004. Genetic programming in classifying large-scale data: an ensemble method. Information Sciences, 163, 85-101.

Figure 1: **Tree-structured representation and genetic operators**

This figure illustrates the tree-structured individuals and the genetic operators of crossover and mutation. The parents individuals in Panel A and B are combined by the crossover operator, and the resultant offspring individuals are shown in Panel C and D. The offspring individual in Panel E (F) is produced by the mutation operator from the individual in Panel A (B).

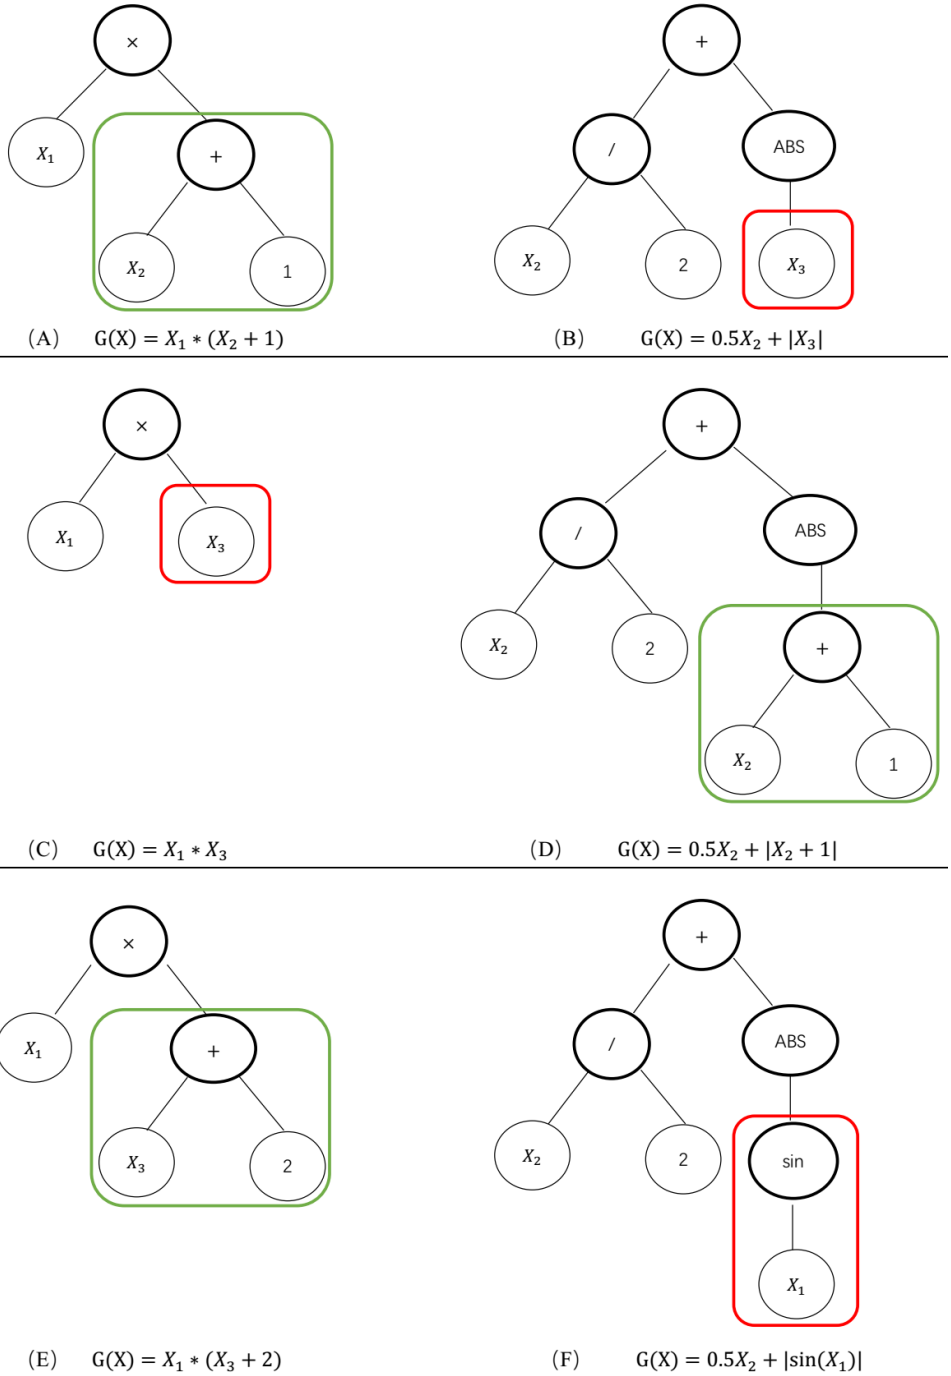


Table 1
Empirical distribution of GP portfolios' Sharpe ratio

This table reports the empirical distribution statistics of the annualized Sharpe ratio of the GP portfolio constructed by Eq. (2) and Eq. (5), and the stock universe is the 100 largest stocks from the S&P 500 index constituents. We re-estimate the GP models for 100 RUNs, producing 100 GP portfolios accordingly. This table reports the empirical distribution statistics (mean, median, 1-, 5-, 95-, and 99-percentile) of the 100 portfolios. We report the results for various Generation parameters of 10, 20, and 40, and the Population is set to 200. The training period is from 1945:01 to 1980:12. The validation period is from 1981:01 to 1990:12. The OOS period is from 1991:01 to 2021:12.

	Average	Median	p1	p5	p95	p99
Panel A: Training period						
Gen=10	1.4089	1.4058	1.2864	1.3391	1.5127	1.5403
Gen=20	1.6017	1.5990	1.4438	1.4981	1.7000	1.7696
Gen=40	1.7933	1.7888	1.5970	1.6841	1.9228	1.9287
Panel B: Validation period						
Gen=10	0.8274	0.8219	0.6695	0.7074	0.9765	1.0567
Gen=20	0.8563	0.8945	0.5908	0.6304	0.9975	1.0957
Gen=40	0.8712	0.9133	0.5093	0.5748	1.1040	1.1469
Panel C: OOS period						
Gen=10	0.9838	0.9754	0.8425	0.8789	1.1147	1.1633
Gen=20	0.9829	0.9800	0.7574	0.8417	1.1542	1.1896
Gen=40	0.9864	0.9787	0.7266	0.7488	1.1809	1.2137

Table 2
Comparison with the optimal combining strategy of Kan et al. (2022)

This table compares the OOS performances of the baseline GP model with the optimal combining strategy of Kan et al. (2022). The first column reports the annualized Sharpe ratio (SR) in the OOS period. The other columns report the p-values of the one-sided tests on H0: the SR for the model in the corresponding column $<$ the SR for the model in the corresponding row. w^* is the optimal portfolio based on the in-sample estimates of the mean and covariance matrix over the OOS period. The $<Pop, Gen>$ for GP is $< 200, 40 >$. We estimate the GP model for 100 RUNs. GP_{Median} is the model with the median of the validation sample SR. GP_{Top5}, GP_{Top10}, GP_{Top15}, GP_{Top20} is the average of the top 5, 10, 15, and 20 models, where the models are sorted by the validation sample SR. The notations of other models are the same as those in Kan et al. (2022). The stock universe is the largest 100 stocks in S&P 500 index.

		p-values of tests on SR difference								
	SR	GP_{Median}	GP_{Top5}	GP_{Top10}	GP_{Top15}	GP_{Top20}	$\hat{w}_{q,t}$	$\hat{w}_{q,t}^{LW2004}$	$\hat{w}_{q,t}^{LW2017}$	$\hat{w}_{q,t}^{MP}$
w^*	1.8005	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Panel A: Portfolio of GP										
$\boldsymbol{GP_{Median}}$	0.9787	-	0.27	0.17	0.08	0.05	0.99	0.96	0.95	0.96
$\boldsymbol{GP_{Top5}}$	1.0345	0.73	-	0.29	0.10	0.07	0.99	0.97	0.96	0.97
$\boldsymbol{GP_{Top10}}$	1.0526	0.83	0.71	-	0.02	0.01	0.99	0.98	0.97	0.98
$\boldsymbol{GP_{Top15}}$	1.0826	0.92	0.90	0.98	-	0.08	1.00	0.98	0.98	0.98
$\boldsymbol{GP_{Top20}}$	1.1001	0.95	0.93	0.99	0.92	-	1.00	0.99	0.98	0.99
Panel B: Invariant optimal portfolio rules										
$\boldsymbol{\hat{w}_{q,t}}$	0.3985	0.01	0.01	0.01	0.00	0.00	-	0.02	0.02	0.20
$\hat{w}_{p,t}$	-0.2657	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
$\hat{w}_{u,t}$	-0.2589	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
$\hat{w}_{BS,t}$	-0.1174	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Panel C: Rules with shrinkage covariance matrix estimators										
$\boldsymbol{\hat{w}_{q,t}^{LW2004}}$	0.5308	0.04	0.03	0.02	0.02	0.01	0.98	-	0.14	0.52
$\hat{w}_{p,t}^{LW2004}$	-0.2206	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
$\boldsymbol{\hat{w}_{q,t}^{LW2017}}$	0.5591	0.05	0.04	0.03	0.02	0.02	0.98	0.86	-	0.63
$\hat{w}_{p,t}^{LW2017}$	-0.1917	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Panel D: Rules with MacKinlay-Pástor single-factor structure										
$\boldsymbol{\hat{w}_{q,t}^{MP}}$	0.5263	0.04	0.03	0.02	0.02	0.01	0.80	0.48	0.37	-
$\hat{w}_{p,t}^{MP}$	0.5035	0.03	0.02	0.01	0.01	0.01	0.67	0.45	0.40	0.46
Panel E: Rules with no-short-sale constraints										
$\hat{w}_{p,t}^{NS}$	0.4182	0.01	0.01	0.01	0.00	0.00	0.54	0.27	0.22	0.29
Panel F: Other rules from portfolio optimization										
$\hat{w}_{g,t}$	0.5120	0.03	0.02	0.02	0.01	0.01	0.97	0.40	0.29	0.46
$\hat{w}_{g,t}^{NS}$	0.7016	0.14	0.10	0.09	0.07	0.06	0.99	0.96	0.95	0.96
$\hat{w}_{KZ3,t}$	0.4085	0.01	0.01	0.01	0.00	0.00	0.61	0.04	0.03	0.22
Panel G: Nonoptimization rules										
1/N	0.7036	0.14	0.10	0.09	0.07	0.06	0.94	0.86	0.82	0.88
KO_{VT}	0.7804	0.22	0.16	0.14	0.11	0.10	0.99	0.98	0.97	0.99
KO_{RT}	0.7268	0.16	0.11	0.10	0.07	0.07	0.98	0.93	0.90	0.94
KO_{BT}	0.6909	0.13	0.10	0.08	0.06	0.06	0.92	0.82	0.78	0.85

Table 3
List of benchmark portfolios and their abbreviations

The table lists the benchmark portfolios and their abbreviations. "GP" stands for genetic programming, "MAXSER" stands for maximum-Sharpe-ratio estimated and sparse regression, "MV" stands for mean-variance portfolio, and "GMV" stands for global minimum variance portfolio.

Portfolio	Abbreviation
Panel A: Baseline models	
Plug-in MV on Fama and French (1993) 3-factors	Factor
Three-fund portfolio by Kan and Zhou (2007)	KZ
MAXSER by Ao, Li, and Zheng (2019)	MAXSER
GP with no transaction cost	**GP-NoCost**
GP with transaction cost	**GP-WithCost**
Panel B: MV/GMV with different covariance matrix estimates	
MV with sample cov	MV-P
MV with linear shrinkage cov	MV-LS
MV with nonlinear shrinkage cov	MV-NLS
MV with nonlinear shrinkage cov adjusted for factor models	MV-NLSF
GMV with linear shrinkage cov	GMV-LS
GMV with linear nonshrinkage cov	GMV-NLS
Panel C: Models with short-sale constraints	
MV with sample cov and short-sale and cross-validation	MV-P-SSCV
MV with linear shrinkage cov and short-sale and cross-validation	MV-LS-SSCV
MV with nonlinear shrinkage cov and short-sale and cross-validation	MV-NLS-SSCV
GP with no transaction cost and short-sale	**GP-NoCost-SS**
GP with transaction cost and short-sale	**GP-WithCost-SS**
Panel D: MV with L_1-norm constraint and cross-validation	
MV with sample cov and L_1-CV	MV-P-L_1CV
MV with linear shrinkage cov and L_1-CV	MV-LS-L_1CV
MV with nonlinear shrinkage cov and L_1-CV	MV-NLS-L_1CV

Table 4
Comparison with MAXSER of Ao et al. (2019)

This table compares the OOS performances of the baseline GP model with the MAXSER of Ao et al. (2019) and other models listed in Table 3. We report the annualized Sharpe ratio, monthly return and monthly volatility. The stock universe is the S&P 500 constituents.

	Annual SR	Mon Rt (%)	Mon Risk (%)
Panel A: Baseline models			
Index	0.302	0.384	4.400
EW	0.344	0.487	4.900
Factor	0.409	0.531	4.500
KZ	0.331	0.831	8.700
MAXSER	0.591	0.904	5.300
GP-NoCost	**1.089**	**0.420**	**1.338**
Panel B: MV/GMV with different covariance matrix estimates			
MV-P	0.257	2.723	36.700
MV-LS	0.093	0.209	7.800
MV-NLS	0.091	0.168	6.400
MV-NLSF	0.141	0.232	5.700
GMV-LS	0.436	0.315	2.500
GMV-NLS	0.467	0.364	2.700
Panel C: Models with short-sale constraint			
MV-P-SSCV	0.316	0.611	6.700
MV-LS-SSCV	0.157	0.263	5.800
MV-NLS-SSCV	0.139	0.233	5.800
GP-NoCost-SS	**0.637**	**1.298**	**7.059**
Panel D: MV with L_1-norm constraint and cross-validation			
MV-P-L_1CV	0.128	4.700	6.700
MV-LS-L_1CV	-0.053	4.800	5.800
MV-NLS-L_1CV	-0.059	4.800	5.800

Table 5
GP performance in the full stock universe

This table reports the performance of the GP portfolio constructed by Eq. (2) and Eq. (5). This table reports the median statistics of the 100 portfolios. We report the results for various Generation parameters of 10, 20, and 40, and the Population is set to 200. The stock universe includes all domestic common stocks listed on NYSE, AMEX, and Nasdaq stock markets. We re-estimate the GP models for 100 RUNs, producing 100 GP portfolios accordingly. The stock universe includes all domestic common stocks listed on NYSE, AMEX, and Nasdaq stock markets.

	Training	Validation	OOS
Panel A: Annualized SR			
Gen=10	3.369	3.564	2.014
Gen=20	3.656	3.462	2.030
Gen=40	3.860	3.555	1.995
Panel B: Month return (%)			
Gen=10	1.700	1.739	1.485
Gen=20	1.696	1.554	1.371
Gen=40	1.679	1.473	1.371

Table 6
Comparison with the machine-learning-based portfolio of Liu et at. (2023)

The table compares the optimal portfolio proposed in this paper with the spread portfolios generated by machine learning methods in Liu et at. (2023). The first 11 columns report the summary statistics for the decile spread portfolio generated by GP (GP_{Spread}) and other machine learning models, and the results are from Table 1 of Liu et at. (2023). The last column reports the GP portfolio constructed by Eq. (2) and Eq. (5) in this paper. We re-estimate the GP models for 100 RUNs, producing 100 GP portfolios accordingly. The stock universe includes all domestic common stocks listed on NYSE, AMEX, and Nasdaq stock markets.

	GP_{Spread}	Ridge	LASSO	ENET	PCR	PLS	NN1	NN2	NN3	NN4	NN5	GP_{MV}
Panel A: Portfolio performance												
H-L	1.53***	0.95***	1.03***	0.98***	0.84***	0.93***	0.80***	1.10***	1.14***	1.12***	1.02***	1.48***
t-stat	(6.00)	(3.14)	(3.26)	(3.29)	(2.71)	(3.07)	(3.24)	(3.39)	(4.17)	(4.63)	(4.00)	(8.26)
Std dev	4.38	5.18	5.15	5.11	4.98	5.20	4.61	4.86	5.37	4.68	4.30	2.55
SR	1.21	0.64	0.69	0.67	0.59	0.62	0.61	0.79	0.74	0.83	0.82	2.01
Skew	0.83	0.08	-0.11	-0.06	0.15	0.06	-0.36	1.02	0.82	0.21	0.15	1.18
Panel B: Correlation matrix												
GP_{Spread}	1.00	0.55	0.56	0.55	0.51	0.55	0.40	0.53	0.54	0.56	0.56	0.69
Ridge	0.55	1.00	0.97	0.97	0.92	1.00	0.58	0.79	0.78	0.70	0.62	0.48
LASSO	0.56	0.97	1.00	0.98	0.93	0.97	0.60	0.77	0.77	0.69	0.63	0.48
ENET	0.55	0.97	0.98	1.00	0.94	0.98	0.60	0.79	0.76	0.70	0.63	0.47
PCR	0.51	0.92	0.93	0.94	1.00	0.92	0.62	0.78	0.71	0.67	0.64	0.49
PLS	0.55	1.00	0.97	0.98	0.92	1.00	0.58	0.80	0.78	0.70	0.61	0.48
NN1	0.40	0.58	0.60	0.60	0.62	0.58	1.00	0.59	0.42	0.54	0.52	0.35
NN2	0.53	0.79	0.77	0.79	0.78	0.80	0.59	1.00	0.70	0.67	0.51	0.50
NN3	0.54	0.78	0.77	0.76	0.71	0.78	0.42	0.70	1.00	0.62	0.51	0.46
NN4	0.56	0.70	0.69	0.70	0.67	0.70	0.54	0.67	0.62	1.00	0.61	0.50
NN5	0.56	0.62	0.63	0.63	0.64	0.61	0.52	0.51	0.51	0.61	1.00	0.53
GP_{MV}	0.69	0.48	0.48	0.47	0.49	0.48	0.35	0.50	0.46	0.50	0.53	1.00

Table 7
OOS performance under the cost penalty coefficient λ of 1

This table reports the OOS performances for the GP model with λ of 1 to train the model in (15). The "true" unite transaction cost $\tilde{c}_{t+1}$ is set to 0.1%. We report the annualized Sharpe ratio, monthly return and monthly volatility. The stock universe is the S&P 500 constituents.

	Annual SR	Mon Rt (%)	Mon Risk (%)
Panel A: Baseline models			
Index	0.302	0.384	4.400
EW	0.330	0.467	4.900
Factor	0.330	0.429	4.500
KZ	0.160	0.402	8.700
MAXSER	0.483	0.739	5.300
GP-WithCost	**0.794**	**0.309**	**1.347**
Panel B: MV/GMV with different covariance matrix estimates			
MV-P	-0.018	-0.186	35.700
MV-LS	0.011	0.025	7.800
MV-NLS	0.022	0.041	6.400
MV-NLSF	0.054	0.089	5.700
GMV-LS	0.350	0.253	2.500
GMV-NLS	0.398	0.310	2.700
Panel C: Models with short-sale constraint			
MV-P-SSCV	0.174	1.183	6.800
MV-LS-SSCV	0.083	0.490	5.900
MV-NLS-SSCV	0.075	0.435	5.800
GP-WithCost-SS	**0.594**	**1.137**	**6.630**
Panel D: MV with L_1-norm constraint and cross-validation			
MV-P-L_1CV	-0.013	-0.061	4.700
MV-LS-L_1CV	-0.112	-0.538	4.800
MV-NLS-L_1CV	-0.110	-0.528	4.800

Table 8
The effects of transaction cost penalty coefficient λ

This table reports the OOS performances of GP portfolios estimated with different λ in the optimization (15), where the "true" unite transaction cost $\tilde{c}_{t+1}$ is set to 0.1%. We examine two special cases of λ: 0 and 1. For a given λ, we estimate the GP model according to (15). Hence, for λ=1, we estimate the GP model in the training sample assuming the transaction cost is $1 \times \tilde{c}_{t+1}$. Similarly, for λ=0, we estimate the GP model in the training sample assuming there is no transaction cost $(0 \times \tilde{c}_{t+1})$, which tends to produce a more severe overfitting problem. Note that in the OOS test, we use the "true" unite transaction cost $\tilde{c}_{t+1}$ to test the GP models estimated with different λ. The stock universe is the S&P 500 constituents.

	Panel A: Short-sale constraint		Panel B: Short-sale allowed	
λ for model training	0	1	0	1
Annual SR	0.548	0.594	0.751	0.794
Mon Rt (%)	1.117	1.137	0.290	0.309
Mon Vol (%)	7.060	6.630	1.338	1.347
Mon Turn	1.810	1.648	1.300	1.254

Table 9
GP_{MV}-implied expected return: ER_{GP-MV} vs $ER_{GP-Spread}$

Panel A of this table reports the value-weighted decile portfolio return sorted by the GP_{MV}-implied expected return (denoted as ER_{GP-MV}). Suppose $\boldsymbol{w}_{MV}$ is the GP portfolio weight in Eq. (2) and Eq. (5). Based on the Markowitz (1952) optimal weight formula, we derive a GP_{MV}-implied expected return as $ER_{GP-MV} = \hat{\Sigma}\boldsymbol{w}_{MV}$. We use three covariance matrices to proxy $\hat{\Sigma}$, that is the historical covariance matrix based on the daily return in the past 3, 6, and 12 months (i.e., $\hat{\Sigma}_{3Mon}$, $\hat{\Sigma}_{6Mon}$, and $\hat{\Sigma}_{12Mon}$). Panel B and Panel C compare the GP_{MV}-implied expected return (ER_{GP-MV}) with the expected return generated by GP under the objective of maximizing the spread portfolio's Sharpe ratio in liu et al. (2023) (denoted as $ER_{GP-Spread}$). Panel B reports decile portfolios sorted by ER_{GP-MV} controlling for $ER_{GP-Spread}$ in OLS cross-section regressions. In specific, we first conduct a cross-section regression to regress ER_{GP-MV} on $ER_{GP-Spread}$, and then sort stock by the resultant residual. Panel C reports decile portfolios sorted by $ER_{GP-Spread}$ controlling for ER_{GP-MV}. The stock universe includes all domestic common stocks listed on NYSE, AMEX, and Nasdaq stock markets. We re-estimate the GP models for 100 RUNs, producing 100 GP portfolios accordingly.

	Panel A: ER_{GP-MV}, with various $\hat{\Sigma}$			Panel B: ER_{GP-MV}, control for $ER_{GP-Spread}$			Panel C: $ER_{GP-Spread}$, control for ER_{GP-MV}		
	$\hat{\Sigma}_{3Mon}$	$\hat{\Sigma}_{6Mon}$	$\hat{\Sigma}_{12Mon}$	$\hat{\Sigma}_{3Mon}$	$\hat{\Sigma}_{6Mon}$	$\hat{\Sigma}_{12Mon}$	$\hat{\Sigma}_{3Mon}$	$\hat{\Sigma}_{6Mon}$	$\hat{\Sigma}_{12Mon}$
Low	0.14	0.13	0.20	0.39	0.39	0.50	1.13	1.10	0.94
2	0.51	0.44	0.53	0.68	0.70	0.56	1.15	1.13	1.11
3	0.64	0.67	0.56	0.91	0.84	0.81	1.06	1.10	1.16
4	0.79	0.74	0.80	0.86	0.81	0.86	0.97	1.03	1.14
5	0.81	0.81	0.90	0.94	0.99	0.97	0.97	0.93	0.93
6	1.14	0.97	0.88	1.13	0.99	0.90	0.93	1.01	0.91
7	1.09	1.07	1.18	1.12	1.02	1.11	1.00	0.93	1.00
8	1.23	1.29	1.23	1.31	1.26	1.19	0.93	0.95	1.05
9	1.30	1.28	1.14	1.12	1.26	1.25	0.96	1.13	1.06
High	1.49	1.61	1.68	1.29	1.38	1.52	1.13	1.14	1.24
H-L	1.34***	1.48***	1.47***	0.90***	0.98***	1.02***	-0.00	0.03	0.30
t-stat	(4.55)	(3.36)	(3.25)	(3.69)	(3.25)	(3.24)	(-0.01)	(0.12)	(1.15)

Table 10
GP_{MV}-implied expected return: ER_{GP-MV} vs other MLs

This table compares the GP_{MV}-implied expected return with the expected return generated by 11 MLs in Liu et al. (2023). The 11 column denotes the 11 machine learning methods in Liu et al. (2023). Suppose $\boldsymbol{w}_{MV}$ is the GP portfolio weight in Eq. (2) and Eq. (5). Based on the Markowitz (1952) optimal weight formula, we derive a GP_{MV}-implied expected return as $ER_{GP-MV} = \hat{\Sigma}\boldsymbol{w}_{MV}$. We use three covariance matrices to proxy $\hat{\Sigma}$, that is the historical covariance matrix based on the daily return in the past 3, 6, and 12 months (i.e., $\hat{\Sigma}_{3Mon}$, $\hat{\Sigma}_{6Mon}$, and $\hat{\Sigma}_{12Mon}$. Panel A reports the value-weighted decile spread portfolio of the GP_{MV}-implied expected return (ER_{GP-MV}), controlling for the expected return of the model in the corresponding column. Panel B reports the value-weighted decile spread portfolio of the expected return of the model in the corresponding column, controlling for the GP_{MV}-implied expected return (ER_{GP-MV}).

	GP_{Spread}	Ridge	LASSO	ENET	PCR	PLS	NN1	NN2	NN3	NN4	NN5
Panel A: Decile spread portfolio of ER_{GP-MV} with various $\hat{\Sigma}$, controlling for expected return of the model in the column											
$\hat{\Sigma}_{3Mon}$	0.90***	1.01***	0.96***	0.94***	0.92***	0.97***	0.97***	0.83***	0.87***	0.91***	1.08***
	(3.69)	(3.96)	(3.74)	(3.64)	(3.57)	(3.80)	(3.83)	(3.23)	(3.53)	(3.55)	(4.41)
$\hat{\Sigma}_{6Mon}$	0.98***	1.04***	1.08***	1.10***	1.24***	1.02***	1.16***	1.07***	1.12***	1.15***	1.06***
	(3.25)	(3.59)	(3.79)	(3.79)	(4.29)	(3.55)	(4.08)	(3.76)	(4.10)	(4.01)	(3.92)
$\hat{\Sigma}_{12Mon}$	1.02***	0.98***	1.08***	1.02***	1.11***	1.00***	1.07***	0.93***	1.03***	1.22***	0.98***
	(3.24)	(3.22)	(3.51)	(3.31)	(3.74)	(3.28)	(3.53)	(2.98)	(3.58)	(3.98)	(3.34)
Panel B: Decile spread portfolio of the model in the column, controlling for ER_{GP-MV} with various $\hat{\Sigma}$											
$\hat{\Sigma}_{3Mon}$	-0.00	0.26	0.22	0.26	0.19	0.26	0.05	0.46*	0.57**	-0.02	-0.03
	(-0.01)	(1.01)	(0.85)	(1.00)	(0.75)	(1.02)	(0.22)	(1.92)	(2.20)	(-0.10)	(-0.14)
$\hat{\Sigma}_{6Mon}$	0.03	0.29	0.23	0.24	0.05	0.28	0.01	0.37	0.53**	-0.19	-0.21
	(0.12)	(1.05)	(0.87)	(0.90)	(0.25)	(1.03)	(0.04)	(1.59)	(2.16)	(-0.75)	(-0.99)
$\hat{\Sigma}_{12Mon}$	0.30	0.23	0.26	0.22	0.08	0.25	0.13	0.17	0.69***	0.06	-0.11
	(1.15)	(0.90)	(1.01)	(0.86)	(0.38)	(0.97)	(0.59)	(0.71)	(2.66)	(0.26)	(-0.47)