

---

# Stochastic Portfolio Theory: A Machine Learning Perspective

---

**Yves-Laurent Kom Samo**  
Machine Learning Research Group  
Oxford-Man Institute of Quantitative Finance  
University of Oxford  
YLKS@ROBOTS.OX.AC.UK

**Alexander Vervuurt**  
Mathematical Institute  
Oxford-Man Institute of Quantitative Finance  
University of Oxford  
VERVUURT@MATHS.OX.AC.UK

## Abstract

In this paper we propose a novel application of Gaussian processes (GPs) to financial asset allocation. Our approach is deeply rooted in *Stochastic Portfolio Theory* (SPT), a stochastic analysis framework introduced by Robert Fernholz that aims at flexibly analysing the performance of certain investment strategies in stock markets relative to benchmark indices. In particular, SPT has exhibited some investment strategies based on company sizes that, under realistic assumptions, outperform benchmark indices with probability 1 over certain time horizons. Galvanised by this result, we consider the inverse problem that consists of learning (from historical data) an optimal investment strategy based on any given set of trading characteristics, and using a user-specified optimality criterion that may go beyond outperforming a benchmark index. Although this inverse problem is of the utmost interest to investment management practitioners, it can hardly be tackled using the SPT framework. We show that our machine learning approach learns investment strategies that considerably outperform existing SPT strategies in the US stock market.

## 1 INTRODUCTION

Stochastic Portfolio Theory (SPT) is a relatively new stream in financial mathematics, initiated and largely developed by Robert Fernholz [2002]. For surveys of the field, see Fernholz and Karatzas [2009] and Vervuurt [2015]. Among many other things, SPT offers an alternative approach to portfolio selection, taking as its selection criterion to outperform the market index (for instance, the S&P 500 index) with probability one. Investment strategies which achieve this are called *relative arbitrages*, and have been constructed in certain classes of market models. The almost-sure comparison between the performance of

certain portfolios and that of the market is facilitated by Fernholz's 'master equation', a pathwise decomposition of this relative performance which is free from stochastic integrals. The foregoing master equation is the main strength of SPT portfolio selection, as it allows one to circumvent the challenges of explicit model postulation and calibration, as well as the (normative) no-arbitrage assumption, that are encountered in the classical approaches to portfolio optimisation. However, there remain several problems in and limitations to the SPT framework as it stands.

First of all, the task of finding relative arbitrages under reasonable assumptions on the market model is difficult, since it is an inverse problem (this has also been noted by Wong [2015]). Namely, given an investment strategy and market assumptions, one can check whether this strategy is a relative arbitrage (although this quickly becomes very hard for more complicated strategies), but the theory itself does not suggest such strategies. As such, the number of relative arbitrages that have been constructed explicitly remains very small. In a practical setting it would be preferable to invert the problem, and learn investment strategies from data using a user-specified performance criterion. In effect, most established investment managers will likely have a strong view on: i) what performance metric to use to evaluate their strategies, and ii) what values for the chosen metric they regard as being exceptional. The chosen performance metric may depart from the excess return relative to the market index, for instance by adjusting for the risk taken. Similarly, outperforming the market index over a certain time horizon  $[0, T]$  with probability 1 might not be good enough for some practitioners, as investors might pull out following disappointing initial performances, leaving the investment manager unable to realise the long-term optimality. Whence, ideally one should aim at learning from market data what investment strategy is likely to perform exceptionally as per the user's view.

Secondly, several market imperfections are ignored in SPT; most notably, the possibility of bankruptcy is excluded. Since the constructed investment strategies typically invest heavily in small-capitalisation stocks, this poses a strong

limitation on the real-world implementability of these portfolios. However, learning optimal investment strategies from the data copes well with bankruptcies as strategies investing in stocks that eventually fail will naturally be rejected as suboptimal. It also allows for the incorporation of transaction costs, which is theoretically challenging and has not yet been addressed in SPT.

Lastly, the SPT set-up has thus far been developed almost exclusively for investment strategies that are driven only by market capitalisations — there have not yet been any constructions of relative arbitrages driven by other factors. Although this simplification eases theoretical analysis, it is a clear restriction as practitioners do consider many more market characteristics in order to exploit market inefficiencies.

We address all of these issues by adopting a Bayesian non-parametric approach. We consider a broad range of investment strategies driven by a function defined on an arbitrary space of trading characteristics (such as the market capitalisation), on which we place a Gaussian process (GP) prior. For a given strategy, the likelihood of it being ‘exceptional’ is derived from a user-defined performance metric (e.g. excess return to the market index, Sharpe ratio, etc) and values thereof that the practitioner considers ‘exceptional’. We then sample from the posterior of the GP driving the ‘exceptional’ strategy using Monte Carlo Markov Chain (MCMC).

The rest of the paper is structured as follows. In section 2 we provide a background on SPT. In section 3 we present our model, and we illustrate that our approach learns strategies that outperform SPT alternatives in section 4. Finally, we discuss our findings and make suggestions for future research in section 5.

## 2 BACKGROUND

We give a brief introduction to SPT, defining the general class of market models within which its results hold, what the portfolio selection criterion is, and how strategies achieving this criterion are constructed.

### 2.1 THE MODEL

In SPT, the stock capitalisations are modelled as Itô processes.<sup>1</sup> Namely, the dynamics of the  $n$  positive stock capitalisation processes  $X_i(\cdot)$ ,  $i = 1, \dots, n$  are described by the following system of SDEs:

$$dX_i(t) = X_i(t) \left( b_i(t) dt + \sum_{\nu=1}^d \sigma_{i\nu}(t) dW_\nu(t) \right), \quad (1)$$

<sup>1</sup>In the recent work Karatzas and Ruf [2016], it has been shown that this can be weakened to a semimartingale model that even allows for defaults.

for  $t \geq 0$  and  $i = 1, \dots, n$ . Here,  $W_1(\cdot), \dots, W_d(\cdot)$  are independent standard Brownian motions with  $d \geq n$ , and  $X_i(0) > 0$ ,  $i = 1, \dots, n$  are the initial capitalisations. We assume all processes to be defined on a probability space  $(\Omega, \mathcal{F}, \mathbb{P})$ , and adapted to a filtration  $\mathbb{F} = \{\mathcal{F}(t)\}_{0 \leq t < \infty}$  that satisfies the usual conditions and contains the filtration generated by the “driving” Brownian motions. We refer the reader to Karatzas and Shreve [1988] for a reference on stochastic calculus.

The *rates of return*  $b_i(\cdot)$ ,  $i = 1, \dots, n$  and *volatilities*  $\sigma(\cdot) = (\sigma_{i\nu}(\cdot))_{1 \leq i \leq n, 1 \leq \nu \leq d}$ , are some unspecified  $\mathbb{F}$ -progressively measurable processes and are assumed to satisfy the integrability condition

$$\sum_{i=1}^n \int_0^T \left( |b_i(t)| + \sum_{\nu=1}^d (\sigma_{i\nu}(t))^2 \right) dt < \infty, \quad \mathbb{P}\text{-a.s.}, \quad (2)$$

for all  $T \in (0, \infty)$ , and the *non-degeneracy* condition

$$\exists \varepsilon > 0 : \xi^T \sigma(t) \sigma^T(t) \xi \geq \varepsilon \|\xi\|^2, \quad (\text{ND}^\varepsilon)$$

for all  $\xi \in \mathbb{R}^n$  and  $t \geq 0$ ,  $\mathbb{P}$ -almost surely.

### 2.2 RELATIVE ARBITRAGE

In this context, one studies investments in the equity market described by (1) using *portfolios*. These are  $\mathbb{R}^n$ -valued and  $\mathbb{F}$ -progressively measurable processes  $\pi(\cdot) = (\pi_1(\cdot), \dots, \pi_n(\cdot))^T$ , where  $\pi_i(t)$  stands for the proportion of wealth invested in stock  $i$  at time  $t$ .

We restrict ourselves to *long-only* portfolios. These invest solely in the stocks, namely, they take values in the closure  $\overline{\Delta_+^n}$  of the set

$$\Delta_+^n = \left\{ x \in \mathbb{R}^n : x_1 + \dots + x_n = 1, \right. \\ \left. 0 < x_i < 1, i = 1, \dots, n \right\}; \quad (3)$$

in particular, there is no money market. Assuming without loss of generality that the number of outstanding shares of each firm is 1, the corresponding wealth process  $V^\pi(\cdot)$  of an investor implementing  $\pi(\cdot)$  is seen to evolve as follows (we normalise the initial wealth to 1):

$$\frac{dV^\pi(t)}{V^\pi(t)} = \sum_{i=1}^n \pi_i(t) \frac{dX_i(t)}{X_i(t)}, \quad V^\pi(0) = 1. \quad (4)$$

In SPT one measures performance, for the most part, with respect to the market index. This is the wealth process  $V^\mu(\cdot)$  that results from a buy-and-hold portfolio, given by the vector process  $\mu(\cdot) = (\mu_1(\cdot), \dots, \mu_n(\cdot))^T$  of *market weights*

$$\mu_i(t) := \frac{X_i(t)}{X_1(t) + \dots + X_n(t)}. \quad (5)$$

**Definition 1.** Let  $T > 0$ . A strong relative arbitrage with respect to the market over the time-horizon  $[0, T]$  is a portfolio  $\pi(\cdot)$  such that

$$\mathbb{P}(V^\pi(T) > V^\mu(T)) = 1. \quad (6)$$

An equivalent way to express this notion, is to say that the portfolio  $\pi(\cdot)$  strongly outperforms  $\mu(\cdot)$  over the time-horizon  $[0, T]$ .  $\square$

Contrast the SPT approach to portfolio selection with other methods such as mean-variance optimisation (originally introduced by Markowitz [1952]) and expected utility maximisation (see for instance Rogers [2013]), where the optimisation of a certain performance criterion determines the portfolio. In SPT, any portfolio that outperforms the market in the sense of (6) is a relative arbitrage, and the amount by which it outperforms the market is theoretically irrelevant.

In practice, one clearly desires this relative outperformance to be as large as possible. Attempts at optimisation over the class of strategies that satisfy (6) have been made by Fernholz and Karatzas [2010], Fernholz and Karatzas [2011], Ruf [2011], Ruf [2013], and Wong [2015]. However, these results are highly theoretical and very difficult to implement. Our data-driven approach circumvents these theoretical complications by optimising a user-defined criterion over the class of functionally-generated portfolios, which we introduce below.

### 2.3 FUNCTIONALLY-GENERATED PORTFOLIOS

A particular class of portfolios, called *functionally-generated portfolios* (or FGPs for short), was introduced and studied by Fernholz [1999].

Consider a function  $\mathbf{G} \in C^2(U, \mathbb{R}_+)$ , where  $U$  is an open neighbourhood of  $\Delta_+^n$  and such that  $x \mapsto x_i D_i \log \mathbf{G}(x)$  is bounded on  $\Delta_+^n$  for  $i = 1, \dots, n$ . Then  $\mathbf{G}$  is said to be the *generating function of the functionally-generated portfolio*  $\pi(\cdot)$ , given, for  $i = 1, \dots, n$ , by

$$\frac{\pi_i(t)}{\mu_i(t)} = \frac{D_i \mathbf{G}(\mu(t))}{\mathbf{G}(\mu(t))} + 1 - \sum_{j=1}^n \mu_j(t) \frac{D_j \mathbf{G}(\mu(t))}{\mathbf{G}(\mu(t))}. \quad (7)$$

Here, we write  $D_i$  for the partial derivative with respect to the  $i^{\text{th}}$  variable, and we will write  $D_{ij}^2$  for the second partial derivative with respect to the  $i^{\text{th}}$  and  $j^{\text{th}}$  variables. Theorem 3.1 of Fernholz [1999] asserts that the performance of the wealth process corresponding to  $\pi(\cdot)$ , when measured relative to the market, satisfies the  $\mathbb{P}$ -almost sure decomposition (often referred to as ‘‘Fernholz’s master equation’’)

$$\log \left( \frac{V^\pi(T)}{V^\mu(T)} \right) = \log \left( \frac{\mathbf{G}(\mu(T))}{\mathbf{G}(\mu(0))} \right) + \int_0^T \mathbf{g}(t) dt, \quad (8)$$

where the quantity

$$\mathbf{g}(t) := - \sum_{i,j=1}^n \frac{D_{ij}^2 \mathbf{G}(\mu(t))}{2\mathbf{G}(\mu(t))} \mu_i(t) \mu_j(t) \tau_{ij}^\mu(t) \quad (9)$$

is called the *drift process* of the portfolio  $\pi(\cdot)$ . Here, we have written  $\tau_{ij}^\mu(\cdot)$  for the *relative covariances*; denoting by  $e_i$  the  $i^{\text{th}}$  unit vector in  $\mathbb{R}^n$ , these are defined for  $1 \leq i, j \leq n$  as

$$\tau_{ij}^\mu(t) := (\mu(t) - e_i)^T \sigma(t) \sigma^T(t) (\mu(t) - e_j). \quad (10)$$

Under suitable conditions on the market model (1), the left hand side of master equation (8) can be bounded away from zero for sufficiently large  $T > 0$ , thus proving that  $\pi(\cdot)$  is an arbitrage relative to the market over  $[0, T]$ . Several FGPs have been shown to outperform the market this way — see Fernholz [2002], Fernholz et al. [2005], Fernholz and Karatzas [2005], Banner and Fernholz [2008], Fernholz and Karatzas [2009], Picková [2014], and Vervuurt and Karatzas [2015]. In fact, Pal and Wong [2014] prove that any relative arbitrage with respect to the market is necessarily of the form (7), if one restricts  $\pi(\cdot)$  to be a functional of the current market weights only.

Strong [2014] proves a generalisation of (8) for portfolios which are deterministic functions not only of the market capitalisations, but also of other observable quantities. Namely, let  $\mathbf{x}(t) = (\mu(t), F)^T$ , with  $F$  a continuous,  $\mathbb{R}^k$ -valued,  $\mathbb{F}$ -progressively measurable process of finite variation, and let  $\mathcal{H} \in C^{2,1}(\mathbb{R}^n \times \mathbb{R}^k, \mathbb{R}_+)$ . By an application of Theorem 3.1 of Strong [2014], for any portfolio

$$\frac{\pi_i(t)}{\mu_i(t)} = \frac{D_i \mathcal{H}(\mathbf{x}(t))}{\mathcal{H}(\mathbf{x}(t))} + 1 - \sum_{j=1}^n \mu_j(t) \frac{D_j \mathcal{H}(\mathbf{x}(t))}{\mathcal{H}(\mathbf{x}(t))}, \quad (11)$$

for  $i = 1, \dots, n$ , the following master equation holds:

$$\log \left( \frac{V^\pi(T)}{V^\mu(T)} \right) = \log \left( \frac{\mathcal{H}(\mathbf{x}(T))}{\mathcal{H}(\mathbf{x}(0))} \right) + \int_0^T \tilde{\mathbf{g}}(t) dt \quad (12)$$

$$- \int_0^T \sum_{l=1}^k D_{n+l} \log \mathcal{H}(\mathbf{x}(t)) dF_l(t).$$

Here (compare with (8) and (9))

$$\tilde{\mathbf{g}}(t) := - \sum_{i,j=1}^n \frac{D_{ij}^2 \mathcal{H}(\mathbf{x}(t))}{2\mathcal{H}(\mathbf{x}(t))} \mu_i(t) \mu_j(t) \tau_{ij}^\mu(t). \quad (13)$$

Although explicit in its decomposition, the modified master equation (12) has so far not been applied in the literature. It is very difficult and unclear how to postulate in what way such ‘extended generating functions’  $\mathcal{H}$  should depend on market information, and what additional covariates to use. It is thus of interest to develop a methodology that makes suggestions for what functions  $\mathcal{H}$  to use, and extracts from market data which signals are significant.

### 2.4 DIVERSITY-WEIGHTED PORTFOLIOS

One of the most-studied FGPs is the *diversity-weighted portfolio* (DWP) with parameter  $p \in \mathbb{R}$ , defined in (4.4)

of Fernholz et al. [2005] as

$$\pi_i^{(p)}(t) := \frac{(\mu_i(t))^p}{\sum_{j=1}^n (\mu_j(t))^p}, \quad i = 1, \dots, n. \quad (14)$$

In Eq. (4.5) of Fernholz et al. [2005] it was shown that this portfolio is a relative arbitrage with respect to  $\mu(\cdot)$  over  $[0, T]$  for any  $p \in (0, 1)$  and  $T > 2 \log n / (\varepsilon \delta p)$ , under the condition (ND $^\varepsilon$ ), and that of *diversity* (D $^\delta$ ), introduced below. The latter says that no single company’s capitalisation can take up more than a certain proportion of the entire market, which can be observed to hold in real markets;

$$\exists \delta \in (0, 1) : \mathbb{P} \left( \max_{\substack{1 \leq i \leq n \\ t \in [0, T]}} \mu_i(t) < 1 - \delta \right) = 1. \quad (\text{D}^\delta)$$

In Vervuurt and Karatzas [2015], this result was extended to the DWP with *negative* parameter  $p$ , and several variations of this portfolio were shown to outperform the market over sufficiently long time horizons and under suitable market assumptions. A simulation using real market data supported the claim that these portfolios have the potential to outperform the market index, as well as their positive-parameter counterparts. Our results strongly confirm this finding, as well as computing the optimal parameter  $p$  — see section 4.

### 3 SOLVING THE INVERSE PROBLEM

We consider solving the inverse problem of SPT: given some investment objective, how to go about learning a suitable trading strategy from the data? In doing so, we aim for a method that:

1. Learns from a large class of candidate investment strategies to uncover possibly intricate strategies from the data, typically by making use of non-parametric generative models for the generating functions;
2. Leverages additional sources of information beyond market capitalisations to uncover better investment strategies;
3. Works irrespective of the practitioner’s investment objective (e.g. achieving a high Sharpe Ratio, outperforming alternative benchmark indices, etc).

#### 3.1 MODEL SPECIFICATION

Let  $\mathcal{X} \subset \mathbb{R}^d$  be a set of trading characteristics, for some  $d \geq 1$ . We consider long-only portfolios of the form

$$\pi_i^f(t) = \frac{f(\mathbf{x}_i(t))}{\sum_{j=1}^n f(\mathbf{x}_j(t))}, \quad i = 1 \dots, n, \quad (15)$$

for some continuous function  $f : \mathcal{X} \rightarrow \mathbb{R}_+$ .

The idea behind this choice of investment portfolios is grounded in the fact that in practice, an investment manager will often have a predefined set of characteristics that he uses to compare stocks, for instance company size, balance sheet variables, credit ratings, sector, momentum, market vs. book value, return on assets, management team, on-line sentiment, technical indicators, ‘beta’, etc. The investment manager will typically choose trading characteristics so that they are informative enough to unveil market inefficiencies. Moreover, two stocks that have ‘similar’ characteristics will receive ‘similar’ weights.

This approach includes as special cases all functionally-generated portfolios in the SPT framework, and in particular the diversity-, entropy- and equally-weighted, as well as market, portfolios. Our more general setting allows for any set of trading characteristics.

The trading opportunities in our framework are revealed through the time evolving trading characteristics  $\mathbf{x}_i(t)$ , and the investment map  $f$  fully determines how to go about seizing these opportunities. Whence, learning an investment strategy in our framework is equivalent to learning an investment map  $f$ . To do so, we consider two families of functions. Firstly, galvanised by the theoretical results of SPT, we consider the case where  $\mathcal{X} = \mathbb{R}_+$  is the set of market weights, and we take  $f$  to be of the parametric form

$$f : \mu \mapsto \mu^p, \quad (16)$$

for  $p \in \mathbb{R}$ , which corresponds to the diversity-weighted portfolio (DWP, see section 2.4). Secondly, in order to capture more intricate trading patterns, and to allow for a more general set of trading characteristics  $\mathcal{X} \subset \mathbb{R}^d$ , we also consider an alternative non-parametric approach in which we take  $\log f$  to be a path of a mean-zero Gaussian process with covariance function  $k$

$$\log f \sim \mathcal{GP}(0, k(\cdot, \cdot)). \quad (17)$$

To learn ‘good’ investment maps, we need to introduce an optimality criterion that encodes the user’s investment objective. To do so, we consider a performance functional  $\mathcal{P}_{\mathcal{D}}$  that maps the logarithm of a candidate investment map to the historical performance  $\mathcal{P}_{\mathcal{D}}(\log f)$  of the portfolio  $\pi^f(\cdot)$  as in Eq. (15) over some finite time horizon, given historical data  $\mathcal{D}$ . An example performance functional is the annualised Sharpe Ratio, defined as

$$\text{SR}(\pi) = \sqrt{B} \frac{\hat{\mathbb{E}}(\{r(1), \dots, r(T)\})}{\hat{\mathbb{S}}(\{r(1), \dots, r(T)\})}, \quad (18)$$

where  $r(t) = \sum_{i=1}^n r_i(t) \pi_i^f(t)$  is the return of our portfolio between time  $t - 1$  and time  $t$ ,  $r_i(t)$  is the return of the  $i$ -th asset between time  $t - 1$  and time  $t$ ,  $B$  represents the number of units of time in a business year (e.g. 252 if the returns are daily), and  $\hat{\mathbb{E}}$  (resp.  $\hat{\mathbb{S}}$ ) denote the sample

mean (resp. sample standard deviation). Another example of a performance functional is the excess return relative to a benchmark portfolio  $\pi^*$

$$\text{ER}(\pi^f | \pi^*) = \prod_{t=1}^T \left( 1 + \sum_{i=1}^n r_i(t) \pi_i^f(t) \right) - \prod_{t=1}^T \left( 1 + \sum_{i=1}^n r_i(t) \pi_i^*(t) \right). \quad (19)$$

The nature of  $\mathcal{P}_{\mathcal{D}}$  (Sharpe ratio, excess return, etc) depends on the portfolio manager; we impose no theoretical restriction.

In the parametric case (Eq. (16)),  $\mathcal{P}_{\mathcal{D}}(\log f)$  is effectively a function of one single variable  $p$ , and we can easily learn the optimal  $p$  using standard optimisation techniques.

In many cases, however, it might be preferable to reason under uncertainty and be Bayesian. To do so, we express the investment manager’s view as to what is a good performance through a likelihood model  $p(\mathcal{D} | \log f)$ , which we may choose to be a probability distribution on  $\mathcal{P}_{\mathcal{D}}(\log f)$

$$\mathcal{L}(\mathcal{P}_{\mathcal{D}}(\log f)) := p(\mathcal{D} | \log f). \quad (20)$$

This is perhaps the most important step of the learning process. Indeed, the Bayesian methods we will develop in the next section aim at learning investment maps that provide an appropriate trade-off between how *likely* the map is in light of training data, and how *consistent* it is with prior beliefs. This will only lead to a profitable investment map if ‘likely’ maps satisfy the manager’s investment objective in-sample and vice-versa. If one chooses the likelihood model such that likely maps are strategies that lose money, then our learning machines will learn strategies that lose money!

Fortunately, it is very straightforward to express that likely investment maps are the ones that match a desired investment objective. For instance, we may use as likelihood model that, given a candidate investment map  $f$ , the extent to which it is good, or equivalently the extent to which it is ‘likely’ to be the function driving the strategy we are interested in learning, is the same as the extent to which the Sharpe Ratio it generates in-sample comes from a Gamma distribution with mean 2.0 and standard deviation 0.5. The positive support of the Gamma distribution renders functions leading to negative in-sample Sharpe ratios of no interest, while the concentration of the distribution over the Sharpe Ratio around 2.0 reflects both our target performance and some tolerance around it. The choice of mean (2.0) and standard deviation (0.5) of the Gamma reflects the risk appetite of the investment manager, while the vanishing tails properly reflect the fact that too high a performance  $\mathcal{P}_{\mathcal{D}}(\log f)$  would likely raise suspicions and too low a performance would not be good enough.

To complete our Bayesian model specification, in the parametric case we place on  $p$  a uniform prior on  $[-8, 8]$ .

### 3.2 INFERENCE

Throughout the rest of this paper we will use as performance functional the total excess — transaction cost adjusted — return (as defined in (19)) relative to the equally weighted portfolio (EWP), which has constant weights

$$\pi_i^{\text{EWP}}(t) = \frac{1}{n}, \quad i = 1, \dots, n, \quad \forall t \geq 0. \quad (21)$$

over the whole training period

$$\mathcal{P}_{\mathcal{D}}(\log f) = \text{ER}(\pi^f | \text{EWP}). \quad (22)$$

We assume a 10bps transaction cost upon rebalancing (i.e. we incur a cost of 0.1% of the notional for each transaction). It is well known to algorithmic (execution) trading practitioners that a good rule of thumb is to expect to pay 10bps when executing an order whose size is 10% of the average daily traded volume (ADV) on liquid stocks. Whence, this assumption is reasonable so long as the wealth invested in each stock does not exceed 10% ADV. When needed, we use as likelihood model

$$\mathcal{L}(\mathcal{P}_{\mathcal{D}}(\log f)) = \gamma(\mathcal{P}_{\mathcal{D}}(\log f); a, b), \quad (23)$$

where we denote  $\gamma(\cdot; a, b)$  the probability density function of the Gamma distribution with mean  $a$  and standard deviation  $b$ . As previously discussed,  $a$  and  $b$  need not be learned as they reflect the investment manager’s risk appetite. In the experiments of the following section, we use  $a = 7.0$  and  $b = 0.5$ . In other words, we postulate that the ideal investment strategy should be such that, starting with a unit of wealth, the terminal wealth over the training period should be on average 7.0 units of wealth higher than the terminal wealth achieved by the equally weighted portfolio over the same trading horizon — this is purposely greedy.

**Frequentist parametric:** The first method of inference we consider consists of directly learning the optimal parameter of the DWP by maximising  $\mathcal{P}_{\mathcal{D}}(\log f)$  for  $p \in [-8, 8]$ . As a comparison, the typical range of  $p$  considered in the SPT literature is  $[-1, 1]$ . To avoid any issue with local maxima, we proceed with brute force maximisation on the uniform grid with mesh size 0.05.<sup>2</sup>

**Bayesian parametric:** The second method of inference we consider consists of using the *Metropolis-Hastings* algorithm (Hastings [1970]) to sample from the posterior distribution over the exponent  $p$  in the DWP case,

$$p(p | \mathcal{D}) \propto \mathcal{L}(\mathcal{P}_{\mathcal{D}}(p)) \times \mathbb{1}(p \in [-8, 8]), \quad (24)$$

where we have rewritten  $\mathcal{L}(\mathcal{P}_{\mathcal{D}}(\log f))$  as  $\mathcal{L}(\mathcal{P}_{\mathcal{D}}(p))$  to make the dependency in  $p$  explicit. We sample a proposal

<sup>2</sup>This took no longer than a couple of seconds in every experiment that we ran.

update  $p^*$  from a Gaussian centred at the current exponent  $p$  and with standard deviation 0.5. The acceptance probability is easily found to be

$$r = \min \left( 1, \frac{\mathcal{L}(\mathcal{P}_{\mathcal{D}}(p^*))}{\mathcal{L}(\mathcal{P}_{\mathcal{D}}(p))} \mathbb{1}(p^* \in [-8, 8]) \right). \quad (25)$$

We note in particular that so long as  $p$  is initialised within  $[-8, 8]$ , the indicator function in Eq. (24) will not cause problems to the Markov chain. We typically run 10,000 MH iterations and discard the first 5,000 as ‘burn-in’. We use the posterior mean exponent learned on training data to trade in our testing horizon following the corresponding DWP

$$\hat{f}(\mu) = \mu^{\mathbb{E}(p|\mathcal{D})}. \quad (26)$$

**Bayesian non-parametric:** The third method of inference we consider is Bayesian and non-parametric. We place a Gaussian process prior on  $\log f$

$$\log f \sim \mathcal{GP}(0, k(\cdot, \cdot)). \quad (27)$$

Given the sizes of datasets we consider in our experiments (more than 3 million training inputs — 500 assets over a 25-year period), we approximate the latent function over a Cartesian grid. This approximation fits nicely with the quantised nature of financial data. We use as covariance function a separable product of Rational Quadratic (RQ) kernels

$$k(x, y) = k_0^2 \prod_{i=1}^d \left( 1 + \frac{(x[i] - y[i])^2}{2\alpha_i l_i^2} \right)^{-\alpha_i}, \quad (28)$$

where the hyper-parameters  $k_0, l_i, \alpha_i > 0$ , on which we place independent log-normal priors are all to be inferred. We found the RQ kernel to be a better choice than the Gaussian kernel as it allows for ‘varying length scales’. Denoting by  $\mathbf{f}$  the values of the investment map over the input grid, we prefer to work with the equivalent whitened representation

$$\log \mathbf{f} = L\mathbf{X}, \quad \mathbf{X} \sim \mathcal{N}(0, I), \quad (29)$$

where  $I$  is the identity matrix,  $K = [k(\mathbf{x}_i, \mathbf{x}_j)]_{i,j \leq N}$  is the Gram matrix over all  $N$  input points,  $K = UDU^T$  is the Singular Value Decomposition (SVD) of  $K$  and  $L = UD^{\frac{1}{2}}$ . We use a Blocked Gibbs sampler (Geman and Geman [1984]) to sample from the posterior

$$p(\log \mathbf{X}, \log k_0, \{\log l_i, \log \alpha_i\}_{i \leq d} | \mathcal{D}) \propto \mathcal{L}(\mathcal{P}_{\mathcal{D}}(L\mathbf{X})) \times p(\log \mathbf{X}) p(\log k_0) \prod_{i=1}^d p(\log l_i) p(\log \alpha_i), \quad (30)$$

where we have rewritten  $\mathcal{L}(\mathcal{P}_{\mathcal{D}}(\log f))$  as  $\mathcal{L}(\mathcal{P}_{\mathcal{D}}(L\mathbf{X}))$  to emphasise that the likelihood is fully defined by  $\mathbf{f} = L\mathbf{X}$ . The whitened representation has two primary advantages.

First, it is robust to ill conditioning as we may always compute  $L$ , even when  $K$  is singular. Second, it creates a hard link between function values and hyper-parameters, so that updating the latter affects the likelihood  $\mathcal{L}(\mathcal{P}_{\mathcal{D}}(\log f))$ , and therefore directly contributes towards improving the training performance  $\mathcal{P}_{\mathcal{D}}(\log f)$ : we found this to improve mixing of the Markov chain. Our Blocked Gibbs sampler alternates between updating  $\log \mathbf{X}$  conditional on hyper-parameters, and updating the hyper-parameters (and consequently  $L$ ) conditional on  $\log \mathbf{X}$ . For both steps we use the *elliptical slice sampling* algorithm (Murray et al. [2010]). The computational bottleneck of our sampler is the computation of the SVD of  $K$ , which we may do very efficiently by exploiting the separability of our kernel and the grid structure of the input space using standard Kronecker techniques (see for instance Saatchi [2011]).

## 4 EXPERIMENTS

The universe of stocks we consider in our experiments are the constituents of the S&P 500 index, accounting for changes in index constituents. We rebalance our portfolios on a daily basis. At the end of each trading day, we determine our target portfolio for the next day, which is acquired at the open of the next trading day. When the constituents of the index are due to change on day  $t$ , our target portfolio at the end of day  $t-1$  relates to the constituents of the index on day  $t$  (which would indeed be known to the market on day  $t-1$ ). As previously discussed, we assume that each transaction incurs a charge of 0.1% of its notional value. The returns we use account for corporate events such as dividends, defaults, M&A’s, etc. Our data sources are the CRSP and Compustat databases, and we use data from 1 January 1992 to 31 December 2014.

In our first experiment, we aim to illustrate that the approaches we propose in this paper *consistently* and *considerably* outperform SPT alternatives over a wide range of market conditions. We consider learning optimal investment strategies as described in the previous section using 10 consecutive years worth of data and testing on the following 5 years. We begin on 1st January 1992 for the first training dataset, and roll both training and testing datasets by one year, which leads to a total of 9 pairs of training and testing subsets. We compare the equally-weighted portfolio (EWP), the market portfolio, the diversity-weighted portfolio where the exponent  $p$  is learned by maximising the evaluation functional (DWP\*), the diversity-weighted portfolio where the exponent  $p$  is learned with MCMC (DWP), the Gaussian process approach using as trading characteristic the logarithm of the market weights (CAP), and the Gaussian process approach using as trading characteristics both the logarithm of the market weights and the return-on-assets (CAP+ROA). The return-on-assets (ROA) on day  $t$  is defined as the ratio between the last net income reported by the company and last total assets reported by the com-

pany known on day  $t$  — we note that this quantity may not change on a daily basis but this does not affect our analysis. The rationale behind using the ROA as additional characteristics is to capture not only how big a company is, but also how well it performs relative to its size.

Table 1 summarises the average over the 9 scenarios of the yearly in-sample and out-of-sample returns plus-minus two standard errors. It can be seen that all learned strategies do indeed outperform the benchmark (EWP) in-sample and out-of-sample. Moreover, the performance is greatly improved by considering non-parametric models, even when the only characteristic considered is the market weight. Analysing such families of strategies within the SPT framework would simply be mathematically intractable. Finally, it can be seen that adding more trading characteristics does indeed add value. Crucially, the CAP+ROA portfolio *considerably* and *consistently* outperforms the benchmark (EWP), both in-sample and out-of-sample.

Table 1: Results of our first experiment on the consistency of our learning algorithms to varying market conditions. IS RET (resp. OOS RET) are in-sample (resp. out-of-sample) average (over the 9 runs in the experiment) yearly returns in %  $\pm$  two standard errors.

PORTFOLIO	IS RET (%)	OOS RET (%)
MARKET	8.56 $\pm$ 1.62	6.23 $\pm$ 2.07
EWP	10.56 $\pm$ 1.67	8.99 $\pm$ 1.85
DWP*	11.94 $\pm$ 2.01	12.51 $\pm$ 1.12
DWP	11.91 $\pm$ 1.99	12.50 $\pm$ 1.11
CAP	26.54 $\pm$ 2.38	22.05 $\pm$ 2.89
CAP+ROA	<b>56.18<math>\pm</math>7.35</b>	<b>25.14<math>\pm</math>2.58</b>

In our second experiment, we aim to illustrate that our approaches are robust to financial crises. To do so, we train our model using data from 1 January 1992 to 31 December 2005, and test the learned strategy between 1 January 2006 and 31 December 2014, which includes the 2008 financial crisis. We compare the same investment strategies as before. The posterior distribution over the exponent  $p$  in the Bayesian parametric method is illustrated in Figure 1. The learned posterior mean investment maps are illustrated in Figure 3. In Table 2 we provide in-sample and out-of-sample average yearly returns as well as out-of-sample Sharpe ratios. Once again, it can be seen that: i) all learned portfolios do indeed outperform the benchmark (EWP) in-sample and out-of-sample, ii) non-parametric methods outperform parametric methods, and iii) adding the ROA as an additional characteristic does indeed add value. These conclusions hold true not only in absolute terms (returns) but also after adjusting for risk (Sharpe Ratio). A more granular illustration of how our method performs during the 2008 financial crisis can be seen in the time series of total wealth provided in Figure 2. It turns out that the ROA

does not only improve the return out-of-sample, but it also has a ‘stabilising effect’ in that the volatility of the wealth process is considerably reduced.

Finally, it is also worth stressing that the shape of the learned investment map in the two-dimensional case (Figure 3) suggests that the investment strategy uncovered by our Bayesian nonparametric approach can hardly be replicated with a parametric model. Once again, it would be near impossible to derive analytical results pertaining to such a portfolio within the SPT framework.




Figure 1: Posterior distribution of the parameter  $p$  of the diversity-weighted portfolio in our second experiment. The model was trained with market data between 1st January 1992 to 31st December 2005.

Table 2: Results of our second experiment on the robustness of the proposed approaches to financial crises. Returns (RET) are yearly equivalents (in %) of the total returns over the whole testing period. The annualised Sharpe Ratio (SR) is as per Eq. (18). IS (resp. OOS) stands for in-sample (resp. out-of-sample).

PORTFOLIO	IS RET (%)	OOS RET (%)	OOS SR
MARKET	9.60	7.90	0.47
EWP	13.46	9.60	0.51
DWP*	14.62	11.74	0.56
DWP	14.62	11.38	0.55
CAP	16.49	18.01	0.60
CAP+ROA	<b>37.54</b>	<b>18.33</b>	<b>0.72</b>

## 5 CONCLUSION & DISCUSSION

The inverse problem of *stochastic portfolio theory* (SPT) is the following problem: given a user-defined portfolio se-




Figure 2: Time series of out-of-sample wealth processes in our second experiment. Models were trained with market data between 1st January 1992 to 31st December 2005, and tested from 1st January 2006 to 31st December 2014.

lection criterion, how does one go about constructing suitable investment strategies that meet the desired investment objective? This problem is extremely challenging to solve within the SPT framework. We propose the first solution to the inverse SPT problem and we demonstrate empirically that the proposed methods consistently and considerably outperform standard benchmarks, and are robust to financial crises.

Unlike the SPT framework, our methods are based solely on historical data rather than stochastic calculus. This allows us to consider a very broad class of candidate investment strategies that includes all SPT strategies as special cases, but crucially contains many investment strategies that cannot be analysed in the SPT framework. Unlike the SPT framework, which almost exclusively considers outperforming the market portfolio using investment strategies that are solely based on market weights, our proposed approach can cope with virtually any user-defined investment objective and can exploit any arbitrary set of trading characteristics. We empirically demonstrate that this added flexibility allows us to uncover more subtle patterns in financial markets, which results in greatly improved performance.

Although the Gaussian process in our model was approximated to be piecewise constant on a grid, there is no theoretical or practical obstacle in using an alternative approximation such as sparse Gaussian processes (Quiñero Candela and Rasmussen [2005]) or string Gaussian processes (Kom Samo and Roberts [2015b]). Our method may be extended to learn even subtler patterns using the non-stationary general purpose kernels of Kom Samo and Roberts [2015a]. Our work may also be extended to al-





Figure 3: Learned logarithm investment maps of the CAP portfolio (top) and the CAP+ROA portfolio (bottom) in our second experiment. In the case of the CAP portfolio, the credible band corresponds to  $\pm 2$  posterior standard deviations.

low for long-short investment strategies (i.e. strategies that allow short-selling). Finally, it would be interesting to develop an online extension of our work that would capture temporal changes in market dynamics.

### Acknowledgements

Yves-Laurent is a Google Fellow in Machine Learning and would like to acknowledge support from the Oxford-Man Institute of Quantitative Finance. Alexander gratefully acknowledges PhD studentships from the Engineering and Physical Sciences Research Council, Nomura, and the Oxford-Man Institute of Quantitative Finance. Wharton Research Data Services (WRDS) was used in preparing the data for this paper. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third-party suppliers.



## References

- Adrian D. Banner and Daniel Fernholz. Short-term relative arbitrage in volatility-stabilized markets. *Ann. Finance*, 4:445–454, 2008.
- Daniel Fernholz and Ioannis Karatzas. On optimal arbitrage. *Ann. Appl. Probab.*, 20(4):1179–1204, 2010.
- Daniel Fernholz and Ioannis Karatzas. Optimal arbitrage under model uncertainty. *Ann. Appl. Probab.*, 21(6):2191–2225, 2011.
- Robert Fernholz. Portfolio generating functions. *Quantitative Analysis in Financial Markets*, River Edge, NJ. World Scientific, 1999.
- Robert Fernholz. *Stochastic Portfolio Theory*. Springer, 2002.
- Robert Fernholz and Ioannis Karatzas. Relative arbitrage in volatility-stabilized markets. *Ann. Finance*, 1:149–177, 2005.
- Robert Fernholz and Ioannis Karatzas. Stochastic portfolio theory: A survey. In Alain Bensoussan and Qiang Zhang, editors, *Handbook of Numerical Analysis. Vol. XV. Special volume: mathematical modeling and numerical methods in finance*, volume 15 of *Handbook of Numerical Analysis*. Elsevier/North-Holland, Amsterdam, 2009.
- Robert Fernholz, Ioannis Karatzas, and Constantinos Kardaras. Diversity and relative arbitrage in equity markets. *Finance Stoch.*, 9(1):1–27, 2005.
- Stuart Geman and Donald Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741, 1984.
- W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 24:97–109, 1970.
- Ioannis Karatzas and Johannes Ruf. Trading strategies generated by Lyapunov functions. *arXiv preprint arXiv:1603.08245*, 2016.
- Ioannis Karatzas and Steven Shreve. *Brownian Motion and Stochastic Calculus*. Volume 113 in the series Probability and its Applications (New York). Springer-Verlag, New York, 1988.
- Yves-Laurent Kom Samo and Stephen Roberts. Generalized spectral kernels. *arXiv preprint arXiv:1506.02236*, 2015a.
- Yves-Laurent Kom Samo and Stephen Roberts. String Gaussian processes. *arXiv preprint arXiv:1507.06977*, 2015b.
- Harry Markowitz. Portfolio selection. *The journal of finance*, 7(1):77–91, 1952.
- Iain Murray, Ryan Prescott Adams, and David J. C. MacKay. Elliptical slice sampling. *JMLR: W&CP*, 9:541–548, 2010.
- Soumik Pal and Ting-Kam Leonard Wong. The geometry of relative arbitrage. *Mathematics and Financial Economics*, pages 1–31, 2014.
- Radka Picková. Generalized volatility-stabilized processes. *Ann. Finance*, 10(1):101–125, 2014.
- Joaquin Quiñonero Candela and Carl Edward Rasmussen. A unifying view of sparse approximate gaussian process regression. *The Journal of Machine Learning Research*, 6:1939–1959, 2005.
- L. C. G. Rogers. *Optimal Investment*. Springer-Verlag, New York, 2013.
- Johannes Ruf. *Optimal Trading Strategies Under Arbitrage*. PhD thesis, Columbia University, 2011.
- Johannes Ruf. Hedging under arbitrage. *Math. Finance*, 23(2):297–317, 2013.
- Yunus Saatchi. *Scalable Inference for Structured Gaussian Process Models*. PhD thesis, University of Cambridge, 2011.
- Winslow Strong. Generalizations of functionally generated portfolios with applications to statistical arbitrage. *SIAM Journal on Financial Mathematics*, 5(1):472–492, 2014.
- Alexander Vervuurt. Topics in Stochastic Portfolio Theory. *arXiv preprint arXiv:1504.02988*, 2015.
- Alexander Vervuurt and Ioannis Karatzas. Diversity-weighted portfolios with negative parameter. *Annals of Finance*, 11(3-4):411–432, 2015.
- Ting-Kam Leonard Wong. Optimization of relative arbitrage. *Annals of Finance*, 11(3-4):345–382, 2015.