
Portfolio Construction Using Stratified Models^b

Jonathan Tuck^a, Shane Barratt^a and Stephen Boyd^a

Abstract

In this chapter we develop models of asset return mean and covariance that depend on some observable market conditions, and use these to construct a trading policy that depends on these conditions, and the current portfolio holdings. After discretizing the market conditions, we fit Laplacian regularized stratified models for the return mean and covariance. These models have a different mean and covariance for each market condition, but are regularized so that nearby market conditions have similar models. This technique allows us to fit models for market conditions that have not occurred in the training data, by borrowing strength from nearby market conditions for which we do have data. These models are combined with a Markowitz-inspired optimization method to yield a trading policy that is based on market conditions. We illustrate our method on a small universe of 18 ETFs, using three well known and publicly available market variables to construct 1000 market conditions, and show that it performs well out of sample. The method, however, is general, and scales to much larger problems, that presumably would use proprietary data sources and forecasts along with publicly available data.

17.1 Introduction

Trading policy.

We consider the problem of constructing a trading policy that depends on some observable market conditions, as well as the current portfolio holdings. We denote the asset daily returns as $y_t \in \mathbb{R}^n$, for $t = 1, \dots, T$. The observable market conditions are denoted as z_t . We assume these are discrete or categorical, so we have $z_t \in \{1, \dots, K\}$. We denote the portfolio asset weights as $w_t \in \mathbb{R}^n$, with $\mathbf{1}^T w_t = 1$, where $\mathbf{1}$ is the vector with all entries one. The trading policy has the form

$$\mathcal{T} : \{1, \dots, K\} \times \mathbb{R}^n \rightarrow \mathbb{R}^n,$$

where $w_t = \mathcal{T}(z_t, w_{t-1})$, i.e., it maps the current market condition and previous portfolio weights to the current portfolio weights. In this chapter we refer to

^a Stanford University

^b Published in *Machine Learning And Data Sciences For Financial Markets*, Agostino Capponi and Charles-Albert Lehalle© 2022 Cambridge University Press.

z_t as the market conditions, since in our example it is derived from market conditions, but in fact it could be anything known before the portfolio weights are chosen, including proprietary forecasts or other data. Our policy \mathcal{T} is a simple Markowitz-inspired policy, based on a Laplacian regularized stratified model of the asset return mean and covariance; see, e.g., Markowitz (1952); Grinold and Kahn (1999); Boyd et al. (2017).

Laplacian regularized stratified model.

We model the asset returns, conditioned on market conditions, as Gaussian,

$$y \mid z \sim \mathcal{N}(\mu_z, \Sigma_z),$$

with $\mu_z \in \mathbb{R}^n$ and $\Sigma_z \in \mathbf{S}_{++}^n$ (the set of symmetric positive definite $n \times n$ matrices), $z = 1, \dots, K$. This is a stratified model, with stratification feature z . We fit this stratified model, i.e., determine the means μ_1, \dots, μ_K and covariances $\Sigma_1, \dots, \Sigma_K$, by minimizing the negative log-likelihood of historical training data, plus a regularization term that encourages nearby market conditions to have similar means and covariances. This technique allows us to fit models for market conditions which have not occurred in the training data, by borrowing strength from nearby market conditions for which we do have data. Laplacian regularized stratified models are discussed in, e.g., Danaher et al. (2014); Saegusa and Shojaie (2016); Tuck et al. (2019); Tuck et al. (2021); Tuck and Boyd (2021a,b). One advantage of Laplacian regularized stratified models is they are interpretable. They are also auditable: we can easily check if the results are reasonable.

This chapter.

In this chapter we present a single example of developing a trading policy as described above. Our example is small, with a universe of 18 ETFs, and we use market conditions that are publicly available and well known. Given the small universe and our use of widely available market conditions, we cannot expect much in terms of performance, but we will see that the trading algorithm performs well out of sample. Our example is meant only as a simple illustration of the ideas; the techniques we describe can easily scale to a universe of thousands of assets, and use proprietary forecasts in the market conditions. We have made the code for this chapter available online at https://github.com/cvxgrp/lrsm_portfolio.

Outline.

We start by reviewing Laplacian regularized models in §17.2. In §17.3 we describe the data records and dataset we use. In §17.4 we describe the economic conditions with which we will stratify our return and risk models. In §17.5 and §17.6 we describe, fit, and analyze the stratified return and risk models, respectively. In §17.7 we give the details of how our stratified return and risk models are used to create the trading policy \mathcal{T} . We mention a few extensions and variations of the methods in §17.8.

17.1.1 Related work

A number of studies show that the underlying covariances of equities change during different market conditions, such as when the market performs historically well or poorly (a “bull” or “bear” market, respectively), or when there is historically high or low volatility (Erb et al., 1994; Longin and Solnik., 2001; Ang and Bekaert, 2003, 2004; Borland, 2012). Modeling the dynamics of underlying statistical properties of assets is an area of ongoing research. Many model these statistical properties as occurring in hard regimes (i.e., where the statistical properties are the same within a given regime), and utilize methods such as hidden Markov models (Ryden et al., 1998; Hastie et al., 2009; Nystrup et al., 2018) or greedy Gaussian segmentation (Hallac et al., 2019) to model the transitions and breakpoints between the regimes. In contrast, this chapter assumes a hard regime model of our statistical parameters, but our chief assumption is, informally speaking, that similar regimes have similar statistical parameters.

Asset allocation based on changing market conditions is a sensible method for active portfolio management (Ang and Bekaert, 2002; Ang and Timmermann, 2011; Nystrup et al., 2015; Petre, 2015). A popular method is to utilize convex optimization control policies to dynamically allocate assets in a portfolio, where the time-varying statistical properties are modeled as a hidden Markov model (Nystrup et al., 2019).

17.2 Laplacian regularized stratified models

In this section we review Laplacian regularized stratified models, focusing on the specific models we will use; for more detail see Tuck et al. (2021); Tuck and Boyd (2021a). We are given data records of the form $(z, y) \in \{1, \dots, K\} \times \mathbb{R}^n$, where z is the feature over which we stratify, and y is the outcome. We let $\theta \in \Theta$ denote the parameter values in our model. The stratified model consists of a choice of parameter $\theta_z \in \Theta$ for each value of z . In this chapter we will construct two stratified models. One is for return, where $\theta_z \in \Theta = \mathbb{R}^n$ is an estimate or forecast of return in market condition z . The other is for return covariance, where $\theta_z \in \Theta = \mathbf{S}_{++}^n$ is the inverse covariance or precision matrix, and \mathbf{S}_{++}^n denotes the set of symmetric positive definite $n \times n$ matrices. (We use the precision matrix since it is the natural parameter in the exponential family representation of a Gaussian, and renders the fitting problems convex.)

To choose the parameters $\theta_1, \dots, \theta_K$, we minimize

$$\sum_{k=1}^K (\ell_k(\theta_k) + r(\theta_k)) + \mathcal{L}(\theta_1, \dots, \theta_K). \quad (17.1)$$

Here ℓ_k is the loss function, that depends on the training data y_i , for $z_i = k$, typically a negative log-likelihood under our model for the data. The function r is the local regularizer, chosen to improve out of sample performance of the model.

The last term in (17.1) is the Laplacian regularization, which encourages neighboring values of z , under some weighted graph, to have similar parameters.

It is characterized by $W \in \mathbf{S}^K$, a symmetric weight matrix with zero diagonal entries and nonnegative off-diagonal entries. The Laplacian regularization has the form

$$\mathcal{L}(\theta_1, \dots, \theta_K) = \frac{1}{2} \sum_{i,j=1}^K W_{ij} \|\theta_i - \theta_j\|^2,$$

where the norm is the Euclidean or ℓ_2 norm when θ_z is a vector, and the Frobenius norm when θ_z is a matrix. We think of W as defining a weighted graph, with edges associated with positive entries of W , with edge weight W_{ij} . The larger W_{ij} is, the more encouragement we give for θ_i and θ_j to be close.

When the loss and regularizer are convex, the problem (17.1) is convex, and so in principle is tractable (Boyd and Vandenberghe, 2004). The distributed method introduced in Tuck et al. (2021), which exploits the properties that the first two terms in the objective are separable across k , while the last term is separable across the entries of the parameters, can easily solve very large instances of the problem.

A Laplacian regularized stratified model typically includes several hyper-parameters, for example that scale the local regularization, or scale some of the entries in W . We adjust these hyper-parameters by choosing some values, fitting the Laplacian regularized stratified model for each choice of the hyper-parameters, and evaluating the true loss function on a (held-out) validation set. (The true loss function is often but not always the same as the loss function used in the fitting objective (17.1).) We choose hyper-parameters that give the least, or nearly least, true loss on the validation data, biasing our choice toward larger values, i.e., more regularization.

We make a few observations about Laplacian regularized stratified models. First, they are interpretable, and we can check them for reasonableness by examining the values θ_z , and how they vary with z . At the very least, we can examine the largest and smallest values of each entry (or some function) of θ_z over $z \in \{1, \dots, K\}$.

Second, we note that a Laplacian regularized stratified model can be created even when we have no training data for some, or even many, values of z . The parameter values for those values of z are obtained by borrowing strength from their neighbors for which we do have data. In fact, the parameter values for values of z for which we have no data are weighted averages of their neighbors. This implies a number of interesting properties, such as a maximum principle: Any such value lies between the minimum and maximum values of the parameter over those values of z for which we have data.

17.3 Dataset

Our example considers $n = 18$ ETFs as the universe of assets, listed in table 17.1. These ETFs were chosen because they broadly represent the market. Each data record has the form (y_t, z_t) , where $y_t \in \mathbb{R}^{18}$ is the daily *active* return of each asset with respect to VTI, an ETF which broadly tracks the total stock market, from

Asset	Description
AGG	iShares Core US Aggregate Bond ETF
DBC	PowerShares DB Commodity Index Tracking Fund
GLD	SPDR Gold Shares
IBB	iShares Nasdaq Biotechnology ETF
ITA	iShares US Aerospace & Defense ETF
PBJ	Invesco Dynamic Food & Beverage ETF
TLT	iShares 20 Plus Year Treasury Bond ETF
VNQ	Vanguard Real Estate Index Fund ETF
VTI	Vanguard Total Stock Market Index Fund ETF
XLB	Materials Select Sector SPDR Fund
XLE	Energy Select Sector SPDR Fund
XLF	Financial Select Sector SPDR Fund
XLI	Industrial Select Sector SPDR Fund
XLK	Technology Select Sector SPDR Fund
XLP	Consumer Staples Select Sector SPDR Fund
XLU	Utilities Select Sector SPDR Fund
XLV	Health Care Select Sector SPDR Fund
XLY	Consumer Discretionary Select Sector SPDR Fund

Table 17.1 Universe of 18 ETFs.

market close on day $t - 1$ until market close on day t , and z_t represents the market condition known by the market close on day $t - 1$, described later in §17.4. (The daily active return of each asset with respect to VTI is the daily return of that asset minus the daily return of VTI.) Henceforth, when we refer to return or risk we mean active return or active risk, with respect to our benchmark VTI. The benchmark VTI has zero active return and risk.

Our dataset spans March 2006 to December 2019, for a total of 3461 data points. We first partition it into two subsets. The first, using data from 2006–2014, is used to fit the return and risk models as well as to choose the hyper-parameters in the return and risk models and the trading policy. The second subset, with data in 2015–2019, is used to test the trading policy. We then randomly partition the first subset into two parts: a training set consisting of 80% of the data records, and a validation set consisting of 20% of the data records. Thus we have three datasets: a training data set with 1779 data points in the date range 2006–2014, a validation set with 445 data points also in the date range 2006–2014, and a test dataset with 1237 data points in the date range 2015–2019. We use 9 years of data to fit our models and choose hyper-parameters, and 5 years of later data to test the trading policy. In order to minimize the influence of outliers in the models, return data in the training and validation datasets were winsorized (clipped) at their 1st and 99th percentiles. The return data in the test dataset was not winsorized.

	Volatility	Inflation	Mortgage
Volatility	1	-0.13	-0.28
Inflation	-	1	0.28
Mortgage	-	-	1

Table 17.2 Correlation of the market indicators over the training and validation period, 2006–2014.

17.4 Stratified market conditions

Each data record also includes the market condition z known on the previous day’s market close. To construct the market condition z , we start with three (real-valued) market indicators.

Market implied volatility.

The volatility of the market is a commonly used economic indicator, with extreme values associated with market turbulence (French et al., 1987; Schwert, 1989; Aggarwal et al., 1999; Chun et al., 2020). Here, volatility is measured by the 15-day moving average of the CBOE volatility index (VIX) on the S&P 500 (Exchange, 2020), lagged by an additional day.

Inflation rate.

The inflation rate measures the percentage change of purchasing power in the economy (Wynne and Sigalla, 1994; Boyd et al., 1996, 2001; Boyd and Champ, 2003; Hung, 2003; Mahyar, 2017). The inflation rate is published by the United States Bureau of Labor Statistics (of Labor Statistics, 2020) as the percent change of the consumer price index (CPI), which measures changes in the price level of a representative basket of consumer goods and services, and is updated monthly.

30-year US mortgage rates.

This metric is the interest rate charged by a mortgage lender on 30-year mortgages, and the change of this rate is an economic indicator correlated with economic spending (Cava, 2016; Sutton et al., 2017). The 30-year US mortgage rate are published by the Federal Home Loan Mortgage Corporation, a public government-sponsored enterprise, and is generally updated weekly (Federal Reserve Economic Data, 2020). Here, this market condition is the 8-week rolling percent change of the 30-year US mortgage rate.

These three economic indicators are not particularly correlated over the training and validation period, as can be seen in table 17.2.

Discretization.

Each of these market indicators is binned into deciles, labeled $1, \dots, 10$. (The decile boundaries are computed using the data up to 2015.) The total number of stratification feature values is then $K = 10 \times 10 \times 10 = 1000$. We can think of z as a 3-tuple of deciles, in $\{1, \dots, 10\}^3$, or encoded as a single value $z \in \{1, \dots, 1000\}$.

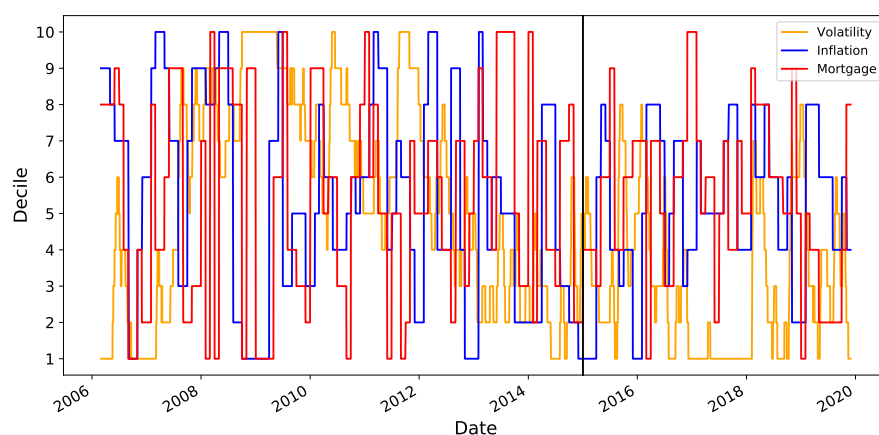


Figure 17.1 Stratification feature values over time. The vertical line at 2015 separates the training and validation period (2006–2014) from the test period (2015–2019).

The market conditions over the entire dataset are shown in figure 17.1, with the vertical line at 2015 indicating the boundary between the training and validation period (2006–2014) and the test period (2015–2019). The average value of $\|z_{t+1} - z_t\|_1$ (interpreting them as vectors in $\{1, \dots, 10\}^3$) is around 0.35, meaning that on each day, the market conditions change by around 0.35 deciles on average.

Data scarcity.

The market conditions can take on $K = 1000$ possible values. In the training/validation datasets, only 346 of 1000 market conditions appear, so there are 654 market conditions for which there are no data points. The most populated market condition, which corresponds to market conditions (9, 0, 0), contains 42 data points. The average number of data points per market condition in the training/validation data is 2.22.

For about 65% of the market conditions, we have *no* training data. This scarcity of data means that the Laplacian regularization is critical in constructing models of the return and risk that depend on the market conditions.

In the test dataset, only 188 of the economic conditions appear. The average number of data points per market condition in the test dataset is 1.24. Only 71 economic conditions appear in both the training/validation and test datasets. In the test data, there are only 442 days (about 36% of the 1237 test data days) in which the market conditions for that day were observed in the training/validation datasets.

Regularization graph.

Laplacian regularization requires a weighted graph that tells us which market conditions are ‘close’. Our graph is the Cartesian product of three chain graphs (Tuck et al., 2021), which link each decile of each indicator to its successor (and predecessor). This graph on the 1000 values of z has 2700 edges. Each edge

connects two adjacent deciles of one of our three economic indicators. We assign three different positive weights to the edges, depending on which indicator they link. We denote these as

$$\gamma_{\text{vol}}, \quad \gamma_{\text{inf}}, \quad \gamma_{\text{mort}}. \quad (17.2)$$

These are hyper-parameters in our Laplacian regularization. Each of the nonzero entries in the weight matrix W is one of these values. For example, the edge between $(3, 1, 4)$ and $(3, 2, 4)$, which connects two values of z that differ by one decile in Inflation, has weight γ_{inf} .

17.5 Stratified return model

In this section we describe the stratified return model. The model consists of a return vector $\theta_z = \mu_z \in \mathbb{R}^{18}$ for each of $K = 1000$ different market conditions, for a total of $Kn = 18000$ parameters.

The loss in (17.1) is a Huber penalty,

$$\ell_k(\mu_k) = \sum_{t:z_t=k} \mathbf{1}^T H(\mu_k - y_t),$$

where H is the Huber penalty (applied entrywise above),

$$H(z) = \begin{cases} z^2, & |z| \leq M \\ 2M|z| - M^2, & |z| > M, \end{cases}$$

where $M > 0$ is the half-width, which we fix at the reasonable value $M = 0.01$. (This corresponds to the 79th percentile of absolute return on the training dataset.) The Huber loss is utilized because it is robust (or less sensitive) to outliers. We use quadratic or ℓ_2 squared local regularization in (17.1),

$$r(\mu_k) = \gamma_{\text{ret,loc}} \|\mu_k\|_2^2,$$

where the positive regularization weight $\gamma_{\text{ret,loc}}$ is another hyper-parameter.

The Laplacian regularization contains the three hyper-parameters (17.2), so overall our stratified return model has four hyper-parameters.

17.5.1 Hyper-parameter search

To choose the hyper-parameters for the stratified return model, we start with a coarse grid search, which evaluates combinations of hyper-parameters over a large range. We evaluate all combinations of

$$\begin{aligned} \gamma_{\text{ret,loc}} &= 0.001, 0.01, 0.1, \\ \gamma_{\text{vol}} &= 1, 10, 100, 1000, 10000, 100000 \\ \gamma_{\text{inf}} &= 1, 10, 100, 1000, 10000, 100000 \\ \gamma_{\text{mort}} &= 1, 10, 100, 1000, 10000, 100000 \end{aligned}$$

Model	Train correlation	Validation correlation
Stratified return model	0.093	0.054
Common return model	0.018	0.001

Table 17.3 Correlations to the true returns over the training set and the held-out validation set for the return models.

a total of 648 combinations, and select the hyper-parameter combination that yields the largest correlation between the return estimates and the returns over the validation set. (Thus, our true loss is negative correlation of forecast and realized returns.) The hyper-parameters

$$(\gamma_{\text{ret,loc}}, \gamma_{\text{vol}}, \gamma_{\text{inf}}, \gamma_{\text{mort}}) = (0.01, 10, 100, 10000)$$

gave the best results over this coarse hyper-parameter grid search.

We then perform a second hyper-parameter grid search on a finer grid of values centered around the best values from the coarse search. We test all combinations of

$$\begin{aligned} \gamma_{\text{ret,loc}} &= 0.0075, 0.01, 0.0125, \\ \gamma_{\text{vol}} &= 2, 5, 10, 20, 50, \\ \gamma_{\text{inf}} &= 20, 50, 100, 200, 500, \\ \gamma_{\text{mort}} &= 2000, 5000, 10000, 20000, 50000, \end{aligned}$$

a total of 375 combinations. The final hyper-parameter values are

$$(\gamma_{\text{ret,loc}}, \gamma_{\text{vol}}, \gamma_{\text{inf}}, \gamma_{\text{mort}}) = (0.01, 20, 50, 5000). \quad (17.3)$$

These can be roughly interpreted as follows. The large value for γ_{mort} tells us that our return model should not vary much with mortgage rate, and the smaller values for γ_{vol} and γ_{inf} tells us that our return model can vary more with volatility and inflation.

17.5.2 Final stratified return model

Table 17.3 shows the correlation coefficient of the return estimates to the true returns over the training and validation sets, for the stratified return model and the common return model, i.e., the empirical mean over the training set. The stratified return model estimates have a larger correlation with the realized returns in both the training set and the validation set.

Table 17.4 summarizes some of the statistics of our stratified return model over the 1000 market conditions, along with the common model value. We can see that each forecast varies considerably across the market conditions. Note that the common model values are the averages over the training data; the median, minimum, and maximum are over the 1000 market conditions.

Asset	Common	Median	Min	Max
AGG	-0.015	-0.064	-0.109	0.045
DBC	-0.049	-0.050	-0.131	0.076
GLD	-0.007	-0.017	-0.111	0.130
IBB	0.040	0.045	-0.053	0.132
ITA	0.022	0.029	-0.062	0.059
PBJ	0.009	0.007	-0.038	0.096
TLT	0.011	-0.053	-0.162	0.092
VNQ	0.015	0.008	-0.229	0.064
VTI	0	0	0	0
XLB	0.003	0.014	-0.033	0.066
XLE	-0.001	0.020	-0.081	0.113
XLF	-0.023	-0.047	-0.341	0.039
XLI	0.008	0.015	-0.053	0.052
XLK	0.001	0.003	-0.045	0.081
XLP	0.006	-0.001	-0.040	0.062
XLU	-0.009	-0.017	-0.067	0.072
XLV	0.012	0.011	-0.029	0.055
XLY	0.014	0.007	-0.048	0.049

Table 17.4 Return predictions, in percent daily return. The first column gives the common return model; the second, third, and fourth columns give median, minimum, and maximum return predictions over the 1000 market conditions for the Laplacian regularized stratified model. All returns are relative to VTI, which has zero return.

17.6 Stratified risk model

In this section we describe the stratified risk model, i.e., a return covariance that depends on z . For determining the risk model, we can safely ignore the (small) mean return, and assume that y_t has zero mean. (The return is small, so the squared return is negligible.) The model consists of $K = 1000$ inverse covariance matrices $\Sigma_k^{-1} = \theta_k \in \mathbf{S}_{+++}^{18}$, indexed by the market conditions. Our stratified risk model has $Kn(n+1)/2 = 171000$ parameters.

The loss in (17.1) is the negative log-likelihood on the training set (scaled, with constant terms ignored),

$$\ell_k(\theta_k) = \text{Tr}(S_k \Sigma_k^{-1}) - \log \det(\Sigma_k^{-1})$$

where $S_k = \frac{1}{n_k} \sum_{t:z_t=k} y_t y_t^T$ is the empirical covariance matrix of the data y for which $z = k$, and n_k is the number of data samples with $z = k$. (When $n_k = 0$, we take $S_k = 0$.) We found that local regularization did not improve the model performance, so we take local regularization $r = 0$. All together our stratified risk model has the three Laplacian hyper-parameters (17.2).

17.6.1 Hyper-parameter search

We start with a coarse grid search over all 216 combinations of

$$\gamma_{\text{vol}} = 0.01, 0.1, 1, 10, 100, 1000,$$

$$\gamma_{\text{inf}} = 0.01, 0.1, 1, 10, 100, 1000,$$

$$\gamma_{\text{mort}} = 0.01, 0.1, 1, 10, 100, 1000,$$

selecting the hyper-parameter combination with the smallest negative log-likelihood (our true loss) on the validation set. The hyper-parameters

$$(\gamma_{\text{vol}}, \gamma_{\text{inf}}, \gamma_{\text{mort}}) = (0.1, 10, 100)$$

gave the best results.

We then perform a second search on a finer grid, focusing on hyper-parameter value near the best values from the coarse search. We evaluate all 125 combinations of

$$\gamma_{\text{vol}} = 0.02, 0.05, 0.1, 0.2, 0.5,$$

$$\gamma_{\text{inf}} = 2, 5, 10, 20, 50,$$

$$\gamma_{\text{mort}} = 20, 50, 100, 200, 500.$$

For the stratified risk model, the final hyper-parameter values chosen are

$$(\gamma_{\text{vol}}, \gamma_{\text{inf}}, \gamma_{\text{mort}}) = (0.2, 20, 50).$$

It is interesting to compare these to the hyper-parameter values chosen for the stratified return model, given in (17.3). Since the losses for return and risk models are different, we can scale the hyper-parameters in the return and risk to compare them. We can see that they are not the same, but not too different, either; both choose γ_{inf} larger than γ_{vol} , and γ_{mort} quite a bit larger than γ_{vol} .

Model	Train loss	Validation loss
Stratified risk model	-6.69	-1.45
Common risk model	3.47	4.99

Table 17.5 Average negative log-likelihood (scaled, with constant terms ignored) over the training and validation sets for the stratified and common risk models.

Asset	Common	Median	Min	Max
AGG	1.314	0.906	0.586	4.135
DBC	1.285	1.070	0.778	3.870
GLD	1.671	1.269	0.982	5.201
IBB	0.905	0.823	0.694	2.120
ITA	0.618	0.557	0.492	1.428
PBJ	0.650	0.513	0.437	1.915
TLT	1.816	1.334	0.809	5.828
VNQ	1.328	0.786	0.666	4.409
VTI	0	0	0	0
XLB	0.771	0.641	0.507	1.703
XLE	1.019	0.857	0.686	2.401
XLF	1.190	0.617	0.389	4.401
XLI	0.500	0.440	0.370	1.045
XLK	0.515	0.465	0.387	1.057
XLP	0.759	0.576	0.455	2.425
XLU	0.882	0.749	0.639	2.186
XLV	0.701	0.509	0.428	2.108
XLY	0.535	0.442	0.355	1.154

Table 17.6 Forecasts of volatility, expressed in percent daily return. The first column gives the common model; the second, third, and fourth columns give median, minimum, and maximum volatility predictions over the 1000 market conditions for the Laplacian regularized stratified model. Volatilities are of return relative to VTI, so VTI has zero volatility.

17.6.2 Final stratified risk model

Table 17.5 shows the average negative log likelihood (scaled, with constant terms ignored) over the training and held-out validation sets, for both the stratified risk model and the common risk model, i.e., the empirical covariance. We can see that the stratified risk model has substantially better loss on the training and validation sets.

Table 17.6 summarizes some of the statistics of our stratified return model asset volatilities, i.e., $(\Sigma_z)_{ii}^{1/2}$, expressed as daily percentages, over the 1000 market conditions, along with the common model asset volatilities. We can see that the predictions vary considerably across the market conditions, with a few varying by a factor almost up to ten. Table 17.7 summarizes the same statistics for the correlation of each asset with AGG, an aggregate bond market ETF. Here we see dramatic variation, for example, the correlation between XLI (an industrials ETF) and AGG varies from -79% to +82% over the market conditions.

Asset	Common	Median	Min	Max
AGG	1	1	1	1
DBC	0.492	0.416	-0.384	0.952
GLD	0.684	0.524	0.093	0.971
IBB	0.250	0.063	-0.585	0.917
ITA	0.024	-0.051	-0.807	0.875
PBJ	0.565	0.384	0.006	0.946
TLT	0.935	0.897	0.803	0.994
VNQ	-0.345	0.021	-0.932	0.652
XLB	-0.214	-0.232	-0.749	0.808
XLE	-0.205	-0.185	-0.935	0.619
XLF	-0.520	-0.289	-0.970	0.042
XLI	-0.107	-0.108	-0.789	0.816
XLK	0.154	0.075	-0.705	0.846
XLP	0.714	0.579	0.344	0.973
XLU	0.555	0.458	0.142	0.939
XLV	0.607	0.429	-0.106	0.962
XLY	-0.061	-0.026	-0.701	0.844

Table 17.7 Forecasts of correlations with the aggregate bond index AGG. The first column gives the common model; the second, third, and fourth columns give median, minimum, and maximum correlation predictions over the 1000 market conditions for the Laplacian regularized stratified model.

17.7 Trading policy and backtest

17.7.1 Trading policy

In this section we give the details of how we use our stratified return and risk models to construct the trading policy \mathcal{T} .

At the beginning of each day t , we use the previous day's market conditions z_t to allocate our current portfolio according to the weights w_t , computed as the solution of the Markowitz-inspired problem (Boyd et al., 2017)

$$\begin{aligned} & \text{maximize} && \mu_{z_t}^T w - \gamma_{sc} \kappa^T (w)_- - \gamma_{tc} \tau_t^T |w - w_{t-1}| \\ & \text{subject to} && w^T \Sigma_{z_t} w \leq \sigma^2, \quad \mathbf{1}^T w = 1, \\ & && \|w\|_1 \leq L_{\max}, \quad w_{\min} \leq w \leq w_{\max}, \end{aligned} \quad (17.4)$$

with optimization variable $w \in \mathbb{R}^{18}$, where $w_- = \max\{0, -w\}$ (elementwise), and the absolute value is elementwise. We describe each term and constraint below.

- *Return forecast.* The first term in the objective, $\mu_{z_t}^T w$, is the expected return under our forecast mean, which depends on the current market conditions.
- *Shorting cost.* The second term $\gamma_{sc} \kappa^T (w)_-$ is a shorting cost, with $\kappa \in \mathbb{R}_+^{18}$ the vector of shorting cost rates. (For simplicity we take the shorting cost rates as constant.) The positive hyper-parameter γ_{sc} scales the shorting cost term, and is used to control our shorting aversion.
- *Transaction cost.* The third term $\gamma_{tc} \tau_t^T |w - w_{t-1}|$ is a transaction cost, with $\tau_t \in \mathbb{R}_+^{18}$ the vector of transaction cost rates used on day t . We take τ_t as one-half the average bid-ask spread of each asset for the previous 15 trading days (excluding the current day). We summarize the bid-ask spreads of each asset over the training and holdout periods in table 17.8. The positive hyper-parameter γ_{tc} scales the transaction cost term, and is used to control the turnover.
- *Risk limit.* The constraint $w^T \Sigma_{z_t} w \leq \sigma^2$ limits the (daily) risk (under our risk model, which depends on market conditions) to σ , which corresponds to an annualized risk of $\sqrt{250}\sigma$.
- *Leverage limit.* The constraint $\|w\|_1 \leq L_{\max}$ limits the portfolio leverage, or equivalently, it limits the total short position $\mathbf{1}^T (w)_-$ to no more than $(L_{\max} - 1)/2$.
- *Position limits.* The constraint $w_{\min} \leq w \leq w_{\max}$ (interpreted elementwise) limits the individual weights.

Parameters.

Some of the constants in the trading policy (17.4) we simply fix to reasonable values. We fix the shorting cost rate vector to $(0.0005)\mathbf{1}$, i.e., 5 basis points for each asset. We take $\sigma = 0.0045$, which corresponds to an annualized volatility (defined as $\sqrt{250}\sigma$) of around 7.1%. We take $L_{\max} = 2$, which means the total short position cannot exceed one half of the portfolio value. (A portfolio with a leverage of 2 is commonly referred to as a *150/50 portfolio*.) We fix the position limits as $w_{\min} = -0.25\mathbf{1}$ and $w_{\max} = 0.4\mathbf{1}$, meaning we cannot short any asset by more than 0.25 times the portfolio value, and we cannot hold more than 0.4 times the portfolio value of any asset.

Asset	Training/validation period	Holdout period
AGG	0.000298	0.000051
DBC	0.000653	0.000324
GLD	0.000112	0.000048
IBB	0.000418	0.000181
ITA	0.000562	0.000175
PBJ	0.000966	0.000637
TLT	0.000157	0.000048
VNQ	0.000394	0.000066
VTI	0.000204	0.000048
XLB	0.000310	0.000098
XLE	0.000181	0.000077
XLF	0.000359	0.000200
XLI	0.000295	0.000079
XLK	0.000324	0.000093
XLP	0.000298	0.000095
XLU	0.000276	0.000099
XLV	0.000271	0.000070
XLY	0.000334	0.000059

Table 17.8 One-half the mean bid-ask spread of each asset, over the training and validation periods and the holdout period.

Hyper-parameters.

Our trading policy has two hyper-parameters, γ_{sc} and γ_{tc} , which control our aversion to shorting and trading, respectively.

17.7.2 Backtests

Backtests are carried out starting from a portfolio of all VTI and a starting portfolio value of $v = 1$ dollars. On day t , after computing w_t as the solution to (17.4), we compute the value of our portfolio v_t by

$$r_{t,\text{net}} = r_t^T w_t - \kappa^T (w_t)_- - (\tau_t^{\text{sim}})^T |w_t - w_{t-1}|, \quad v_t = v_{t-1}(1 + r_{t,\text{net}}),$$

Here $r_t \in \mathbb{R}^{18}$ is the vector of asset returns on day t , $r_t^T w_t$ is the gross return of the portfolio for day t , τ_t^{sim} is one-half the realized bid-ask spread on day t , and $r_{t,\text{net}}$ is the net return of the portfolio for day t including shorting and transaction costs. In particular, *our backtests take shorting and transaction costs into account*. Note also that in the backtests, we use the actual realized bid-ask spread on that day (which is not known at the beginning of the day) to determine the true transaction cost, whereas in the policy, we use the trailing 15 day average (which is known at the beginning of the day).

Our backtest is a bit simplified. Our simulation assumes dividend reinvestment. We account for the shorting and transaction costs by adjusting the portfolio return, which is equivalent to splitting these costs across the whole portfolio; a more careful treatment might include a small cash account. For portfolios of very high

	Return	Risk
Train	11.9%	6.25%
Validation	10.2%	6.88%

Table 17.9 Annualized return and risk for the stratified model policy over the train and validation periods.

value, we would add an additional nonlinear transaction cost term, for example proportional to the 3/2-power or the square of $|w_t - w_{t-1}|$ (Almgren and Chriss, 2000; Boyd et al., 2017).

17.7.3 Hyper-parameter selection

To choose values of the two hyper-parameters in the trading policy, we carry out multiple backtest simulations over the training set. We evaluate these backtest simulations by their realized return (net, including costs) over the validation set.

We perform a grid search, testing all 625 pairs of 25 values of each hyper-parameter logarithmically spaced from 0.1 to 10. The annualized return on the validation set, as a function of the hyper-parameters, are shown in Figure 17.2. We choose the final values

$$\gamma_{sc} = 8.25, \quad \gamma_{tc} = 1.47,$$

shown on Figure 17.2 as a star.

These values are themselves interesting. Roughly speaking, we should plan our trades as if the shorting cost were more than 8.25 times the actual cost, and the transaction cost is about 1.5 times the true transaction cost. The blue and purple region at the bottom of the heat map indicates poor validation performance when the transaction cost parameter is too low, i.e., the policy trades too much.

Table 17.9 gives the annualized return and risk for the policies over the train and validation periods.

Common model trading policy.

We will compare our stratified model trading policy to a common model trading policy, which uses the constant return and risk models, along with the same Markowitz policy (17.4). In this case, none of the parameters in the optimization problem change with market conditions, and the only parameter that changes in different days is w_{t-1} , the previous day's asset weights, which enters into the transaction cost.

We also perform a grid search for this trading policy, over the same 625 pairs of the hyper-parameters. For the common model trading policy, we choose the final values

$$\gamma_{sc} = 1, \quad \gamma_{tc} = 0.38.$$

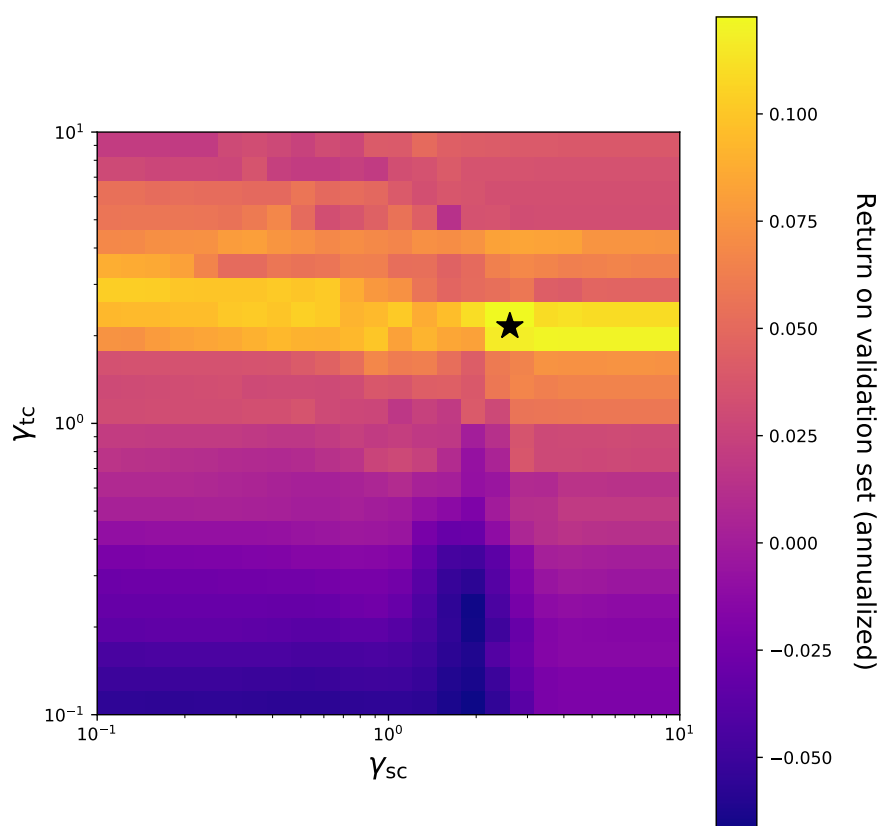


Figure 17.2 Heatmap of the annualized return on the validation set as a function of the two hyper-parameters γ_{sc} and γ_{tc} . The star shows the hyper-parameter combination used in our trading policy.

17.7.4 Final trading policy results

We backtest our trading policy on the test dataset, which includes data from 2015–2019. We remind the reader that no data from this date range was used to create, tune, or validate any of the models, or to choose any hyper-parameters. For comparison, we also give results of a backtest using the constant return and risk models.

Figure 17.3 plots the economic conditions over the test period (top) as well as the active portfolio value (i.e., value above the benchmark VTI) for our stratified model and common model. Buying and holding the benchmark VTI gives zero active return, and a constant active portfolio value of 1. The superior performance of the stratified model policy, e.g., higher Sharpe ratio, is evident in this plot.

Table 17.10 shows the annualized active return, annualized active risk, annualized active Sharpe ratio (return divided by risk), and maximum drawdown of the active portfolio value for the policies over the test period. We remind the reader

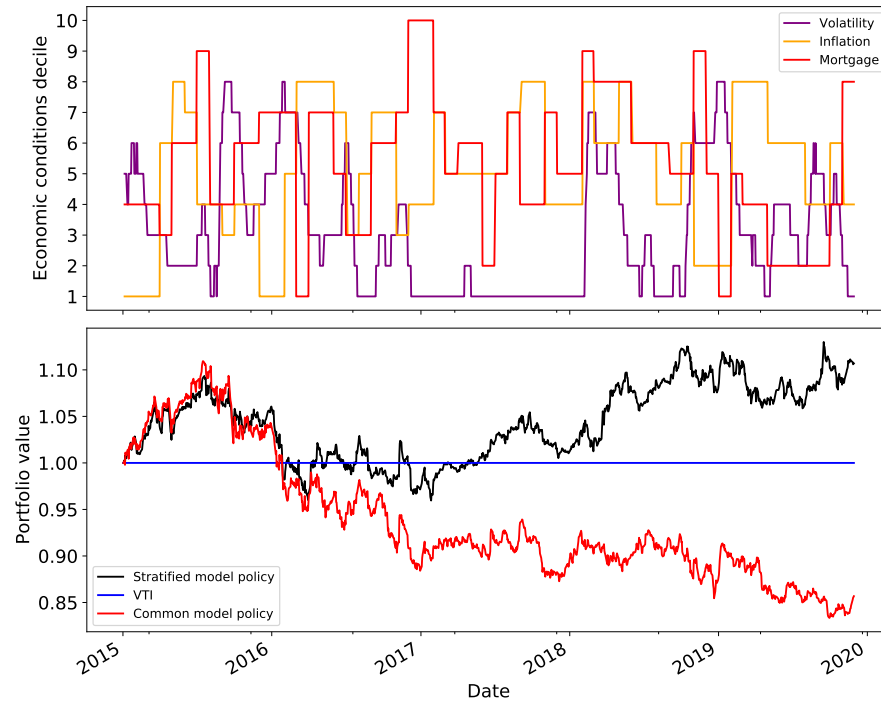


Figure 17.3 Plot of economic conditions (top) and cumulative portfolio value for the stratified model and the common model (bottom) over the test period. The horizontal blue line is the cumulative portfolio value for buying and holding the benchmark VTI.

	Return	Risk	Sharpe ratio	Maximum drawdown
Stratified model policy	2.55%	8.42%	0.302	13.4%
Common model policy	0.003%	7.47%	0.038	16.3%

Table 17.10 Annualized active return, active risk, active Sharpe ratios, and maximum drawdown of the active portfolio value for the three policies over the test period (2015–2019).

that we are fully accounting for the shorting and transaction cost, so the turnover of the policy is accounted for in these backtest metrics.

The results are impressive when viewed in the following light. First, we are using a very small universe of only 18 ETFs. Second, our trading policy uses only three widely available market conditions, and indeed, only their deciles. Third, the policy was entirely developed using data prior to 2015, with no adjustments made for the next five years. (In actual use, one would likely re-train the model periodically, perhaps every quarter or year.)

Comparison of stratified and constant policies.

In Figure 17.4, we plot the asset weights of the stratified model policy (top) and of the common model policy (bottom), over the test period. (The variations in

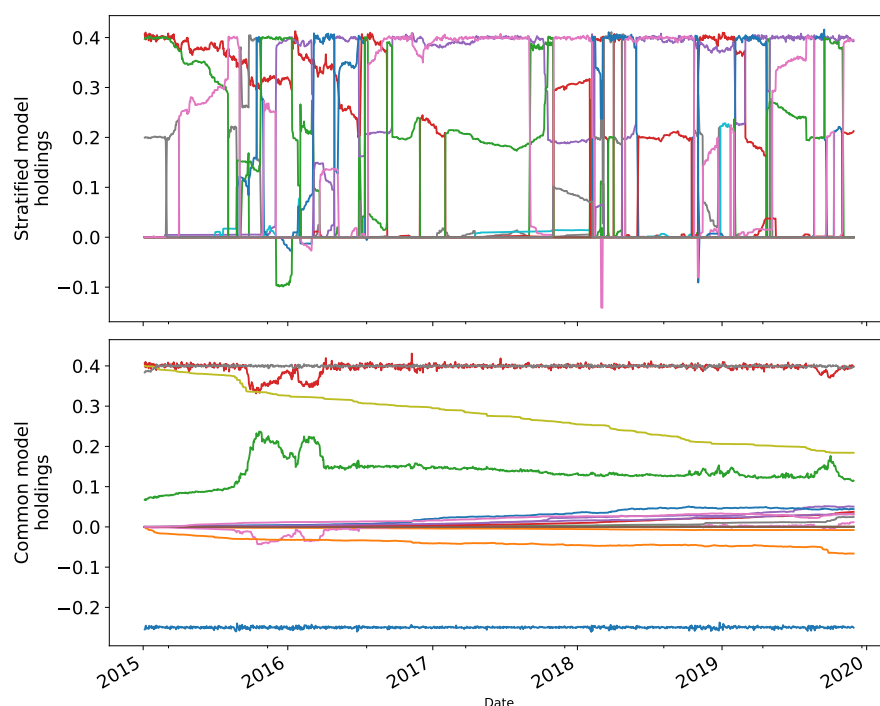


Figure 17.4 Asset weights of the stratified model policy (top) and of the common model policy (bottom), over the test period. The first time period asset weights, which are all VTI, are not shown.

the common model policy holdings come from a combination of a daily rebalancing of the assets and the transaction cost model.) The top plot shows that the weights in the stratified policy change considerably with market conditions. The common model policy is mainly concentrated in just seven assets, GLD (gold), IBB (biotech), ITA (aerospace & defense), XLE (energy), XLV (health care), and XLY (consumer discretionary) (which is effectively cash when considering active returns and risks). Notably, both portfolios are long-only.

Factor analysis.

We fit a linear regression model of the active returns of the two policies over the test set to four of the Fama–French factors (Fama and French, 1992, 1993; French, 2021):

- *MKTRF*, the value-weighted return of United States equities, minus the risk free rate,
- *SMB*, the return on a portfolio of small size stocks minus a portfolio of big size stocks,
- *HML*, the return on a portfolio of value stocks minus a portfolio of growth stocks, and

Factor	Stratified model policy	Common model policy
MKTRF	-0.001362	0.139547
SMB	0.279307	0.235330
HML	-0.361305	-0.448571
UMD	-0.174945	-0.108064
Alpha	0.000085	-0.000215

Table 17.11 The top four rows give the regression model coefficients of the active portfolio returns on the Fama–French factors; the fifth row gives the intercept or alpha value.

- *UMD*, the return on a portfolio of high momentum stocks minus a portfolio of low or negative momentum stocks.

We also include an intercept term, commonly referred to as alpha. Table 17.11 gives the results of these fits. Relative to the common model policy, the stratified model policy active returns are much less positively correlated to the market, shorter the size factor, longer the value factor, and shorter the momentum factor. Its active alpha is around 2.13% annualized. (The common model policy’s active alpha is around -5.38% annualized.) While not very impressive on its own, this alpha seems good considering it was accomplished with just 18 ETFs, and using only three widely available quantities in the policy.

17.8 Extensions and variations

We have presented a simple (but realistic) example only to illustrate the ideas, which can easily be applied in more complex settings, with a far larger universe, a more complex trading policy, and using proprietary forecasts of returns and quantities used to judge market conditions. We describe some extensions and variations on our method below.

Multi-period optimization.

For simplicity we use a policy that is based on solving a single-period Markowitz problem. The entire method immediately extends to policies based on multi-period optimization. For example, we would fit separate stratified models of return and risk for the next 1-day, 5-day, 20-day, and 60-day periods (roughly daily, weekly, monthly, quarterly), all based on the same current market conditions. These data are fed into a multi-period optimizer as described in Boyd et al. (2017).

Joint modeling of return and risk.

In this chapter we have created separate Laplacian regularized stratified models for return and risk. The advantage of this approach is that we can judge each model separately (and with different true objectives), and use different hyperparameter values. It is also possible to fit the return mean and covariance *jointly*, in one stratified model, using the natural parameters in the exponential family for

a Gaussian, Σ^{-1} and $\Sigma^{-1}\mu$. The resulting log-likelihood is jointly concave, and a Laplacian regularized model can be directly fit.

Low-dimensional economic factors.

When just a handful (such as in our example, three) base quantities are used to construct the stratified market conditions, we can bin and grid the values as we do in this chapter. This simple stratification of market conditions preserves interpretability. If we wish to include more raw data in our stratification of market conditions, simple binning and enumeration is not practical. Instead we can use several techniques to handle such situations. The simplest is to perform dimensionality-reduction on the (higher-dimensional) economic conditions, such as principal component analysis (Pearson, 1901) or low-rank forecasting (Barratt et al., 2020), and appropriately bin these low-dimensional economic conditions. These economic conditions may then be related on a graph with edge weights decided by an appropriate method, such as nearest neighbor weights.

Structured covariance estimation.

It is quite common to model the covariance matrix of returns as having structure, e.g., as the sum of a diagonal matrix plus a low-rank matrix (Richard et al., 2012; Fan et al., 2016). This structure can be added by a combination of introducing new variables to the model and encoding constraints in the local regularization. In many cases, this structure constraint turns the stratified risk model fitting problem into a non-convex problem, which may be solved approximately.

Multi-linear interpolation.

In the approach presented above, the economic conditions are categorical, i.e., take on one of $K = 1000$ possible values at each time t , based on the deciles of three quantities. A simple extension is to use multi-linear interpolation (Weiser and Zarantonello, 1988; Davies, 1997) to determine the return and risk to use in the Markowitz optimizer. Thus we would use the actual quantile of the three market quantities, and not just their deciles. In the case of risk, we would apply the interpolation to the precision matrix Σ_r^{-1} , the natural parameter in the exponential family description of a Gaussian.

End-to-end hyper-parameter optimization.

In the example presented in this chapter there are a total of nine hyper-parameters to select. We keep things simple by separately optimizing the hyper-parameters for the stratified return model, the stratified risk model, and the trading policy. This approach allows each step to be checked independently. It is also possible to simultaneously optimize all of the hyper-parameters with respect to a single backtest, using, for example, CVXPYlayers (Agrawal et al., 2019, 2020) to differentiate through the trading policy.

Stratified ensembling.

The methods described in this chapter can be used to combine or ensemble a collection of different return forecasts or signals, whose performance varies with

market (or other) conditions. We start with a collection of return predictions, and combine these (ensemble them) using weights that are a function of the market conditions. We develop a stratified selection of the combining weights.

17.9 Conclusions

We argue that stratified models are interesting and useful in portfolio construction and finance. They can contain a large number of parameters, but unlike, say, neural networks, they are fully interpretable and auditable. They allow arbitrary variation across market conditions, with Laplacian regularization there to help us come up with reasonable models even for market conditions for which we have no training data. The maximum principle mentioned on page 320 tells us that a Laplacian regularized stratified model will never do anything crazy when it encounters values of z that never appeared in the training data. Instead it will use a weighted sum of other values for which we do have training data. These weights are not just any weights, but ones carefully chosen by validation.

The small but realistic example we have presented is only meant to illustrate the ideas. The very same ideas and method can be applied in far more complex and sophisticated settings, with a larger universe of assets, a more complex trading policy, and incorporating proprietary data and forecasts.

Acknowledgements

The authors gratefully acknowledge discussions with and suggestions from Ronald Kahn, Raffaele Savi, and Andrew Ang. We thank an anonymous reviewer for making several suggestions on an early draft, as well as catching some inconsistencies between the data and algorithm as described in the chapter and as implemented in the code. Jonathan Tuck is supported by a Stanford Graduate Fellowship in Science and Engineering.

References

- Aggarwal, R., Inclan, C., and Leal, R. 1999. Volatility in Emerging Stock Markets. *Journal of Financial and Quantitative Analysis*, **34**(1), 33–55.
- Agrawal, A., Amos, B., Barratt, S., Boyd, S., Diamond, S., and Kolter, Z. 2019. Differentiable Convex Optimization Layers. In: *Advances in Neural Information Processing Systems*.
- Agrawal, A., Barratt, S., Boyd, S., and Stellato, B. 2020 (10–11 Jun). Learning Convex Optimization Control Policies. Pages 361–373 of: *Proceedings of the 2nd Conference on Learning for Dynamics and Control*.
- Almgren, R., and Chriss, N. 2000. Optimal execution of portfolio transactions. *Journal of Risk*, **3**(2), 5–39.
- Ang, A., and Bekaert, G. 2002. International Asset Allocation with Regime Shifts. *Review of Financial Studies*, **15**(4), 1137–1187.
- Ang, A., and Bekaert, G. 2003 (Nov.). *How do Regimes Affect Asset Allocation?* Tech. rept. 10080. National Bureau of Economic Research.
- Ang, A., and Bekaert, G. 2004. How Regimes Affect Asset Allocation. *Financial Analysts Journal*, **60**(2), 86–99.

- Ang, A., and Timmermann, A. 2011 (June). *Regime Changes and Financial Markets*. Tech. rept. 17182. National Bureau of Economic Research.
- Barratt, S., Dong, Y., and Boyd, S. 2020. *Low rank forecasting*. ArXiv 2101.12414.
- Borland, L. 2012. Statistical signatures in times of panic: markets as a self-organizing system. *Quantitative Finance*, **12**(9), 1367–1379.
- Boyd, J., and Champ, B. 2003. Inflation and financial market performance: what have we learned in the last ten years? Tech. rept. 0317. Federal Reserve Bank of Cleveland.
- Boyd, J., Levine, R., and Smith, B. 1996 (Oct.). Inflation and financial market performance. Tech. rept. Federal Reserve Bank of Minneapolis.
- Boyd, J., Levine, R., and Smith, B. 2001. The impact of inflation on financial sector performance. *Journal of Monetary Economics*, **47**(2), 221 – 248.
- Boyd, S., and Vandenberghe, L. 2004. *Convex Optimization*. Cambridge University Press.
- Boyd, S., Busseti, E., Diamond, S., Kahn, R., Koh, K., Nystrup, P., and Speth, J. 2017. Multi-Period Trading via Convex Optimization. *Foundations and Trends in Optimization*, **3**(1), 1–76.
- Cava, G. La. 2016 (July). *Housing prices, mortgage interest rates and the rising share of capital income in the United States*. BIS Working Papers 572. Bank for International Settlements.
- Chun, D., Cho, H., and Ryu, D. 2020. Economic indicators and stock market volatility in an emerging economy. *Economic Systems*, **44**(2), 100788.
- Danaher, P., Wang, P., and Witten, D. 2014. The joint graphical lasso for inverse covariance estimation across multiple classes. *Journal of the Royal Statistical Society*, **76**(2), 373–397.
- Davies, S. 1997. Multidimensional Triangulation and Interpolation for Reinforcement Learning. Pages 1005–1011 of: Mozer, M. C., Jordan, M., and Petsche, T. (eds), *Advances in Neural Information Processing Systems 9*. MIT Press.
- Erb, C., Harvey, C., and Viskanta, T. 1994. Forecasting International Equity Correlations. *Financial Analysts Journal*, **50**(6), 32–45.
- Exchange, Chicago Board Options. 2020. CBOE Volatility Index. <http://www.cboe.com/vix>.
- Fama, E., and French, K. 1992. The cross-section of expected stock returns. *Journal of Finance*, **47**(2), 427–465.
- Fama, E., and French, K. 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, **33**(1), 3–56.
- Fan, J., Liao, Y., and Liu, H. 2016. An overview of the estimation of large covariance and precision matrices. *Econometrics Journal*, **19**(1), C1–C32.
- Federal Reserve Economic Data, Federal Reserve Bank of St. Louis. 2020. 30-Year Fixed Rate Mortgage Average in the United States (MORTGAGE30US). <https://fred.stlouisfed.org/series/MORTGAGE30US>.
- French, K. 2021. Description of Fama/French Factors. https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html#Research.
- French, K., Schwert, W., and Stambaugh, R. 1987. Expected stock returns and volatility. *Journal of Financial Economics*, **19**(1), 3.
- Grinold, R., and Kahn, R. 1999. *Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk*. McGraw-Hill.
- Hallac, D., Nystrup, P., and Boyd, S. 2019. Greedy Gaussian Segmentation of Multivariate Time Series. *Advances in Data Analysis and Classification*, **13**(3), 727–751.
- Hastie, T., Tibshirani, R., and Friedman, J. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer.
- Hung, F.-S. 2003. Inflation, financial development, and economic growth. *International Review of Economics & Finance*, **12**(1), 45–67.
- Longin, F., and Solnik, B. 2001. Correlation Structure of International Equity Markets During Extremely Volatile Periods. *Journal of Finance*, **56**(2), 649–676.
- Mahyar, H. 2017. The Effect Of Inflation On Financial Development Indicators In Iran (2000-2015). *Studies in Business and Economics*, **12**(2), 53–62.

- Markowitz, H. 1952. Portfolio selection. *Journal of Finance*, **7**(1), 77–91.
- Nystrup, P., Hansen, B., Madsen, H., and Lindström, E. 2015. Regime-Based Versus Static Asset Allocation: Letting the Data Speak. *Journal of Portfolio Management*, **42**(1), 103–109.
- Nystrup, P., Madsen, H., and Lindström, E. 2018. Dynamic portfolio optimization across hidden market regimes. *Quantitative Finance*, **18**(1), 83–95.
- Nystrup, P., Boyd, S., Lindström, E., and Madsen, H. 2019. Multi-period portfolio selection with drawdown control. *Annals of Operations Research*, **282**(1), 245–271.
- of Labor Statistics, United States Bureau. 2020. Consumer Price Index. <https://www.bls.gov/cpi/>.
- Pearson, K. 1901. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, **2**(11), 559–572.
- Petre, G. 2015. A Case for Dynamic Asset Allocation for Long Term Investors. *Procedia Economics and Finance*, **29**, 41–55.
- Richard, E., Savalle, P.-A., and Vayatis, N. 2012. Estimation of Simultaneously Sparse and Low Rank Matrices. Pages 51–58 of: *Proceedings of the 29th International Conference on Machine Learning*. Madison, WI, USA: Omnipress.
- Ryden, T., Terasvirta, T., and Asbrink, S. 1998. Stylized Facts of Daily Return Series and the Hidden Markov Model. *Journal of Applied Econometrics*, **13**(3), 217–244.
- Saegusa, T., and Shojaie, A. 2016. Joint estimation of precision matrices in heterogeneous populations. *Electronic Journal of Statistics*, **10**(1), 1341–1392.
- Schwert, W. 1989. Why Does Stock Market Volatility Change Over Time? *Journal of Finance*, **44**(5), 1115–1153.
- Sutton, G., Mihaljek, D., and Subelytė, A. 2017 (Oct.). Interest rates and house prices in the United States and around the world. BIS Working Papers 665. Bank for International Settlements.
- Tuck, J., and Boyd, S. 2021a. Fitting Laplacian Regularized Stratified Gaussian Models. *Optimization and Engineering*. <https://doi.org/10.1007/s11081-021-09611-5>.
- Tuck, J., and Boyd, S. 2021b. Eigen-stratified models. *Optimization and Engineering*. <https://doi.org/10.1007/s11081-020-09592-x>.
- Tuck, J., Hallac, D., and Boyd, S. 2019. Distributed Majorization-Minimization for Laplacian Regularized Problems. *IEEE/CAA Journal of Automatica Sinica* **6**(1), 45–52.
- Tuck, J., Barratt, S., and Boyd, S. 2021. A Distributed Method for Fitting Laplacian Regularized Stratified Models. *Journal of Machine Learning Research*. To appear.
- Weiser, A., and Zarantonello, S. 1988. A Note on Piecewise Linear and Multilinear Table Interpolation in Many Dimensions. *Mathematics of Computation*, **50**(181), 189–196.
- Wynne, M., and Sigalla, F. 1994. The consumer price index. *Economic and Financial Policy Review*, **2**(Feb), 1–22.