Portfolio Optimization under Time-Varying Economic Regimes

Cory Turner
ICME
Stanford University
cgturner@stanford.edu

Jiho Han
MS&E
Stanford University
jihohan@stanford.edu

I. Introduction

Our paper aims to answer the question: Can the macroeconomic condition be used to better allocate investments among different sectors of a market? Motivated by classical Markowitz portfolio theory, and prior work in attempting to incorporate exogenous economic factors in portfolio optimization [1][2][3][4], we explore three approaches to dynamically rebalance a portfolio as market or economic regimes change continuously over time. Traditionally, portfolio theory assumes that a set of assets can be characterized by a mean return and a covariance matrix, which are stationary over time. Empirically, it has been found [5] that this stationary assumption does not hold; indeed, the behavior of a group of assets is materially dependent on many observable and unobservable economic factors which change over time.

A. Background: Markowitz Portfolio Theory

Given a vector of returns for k sectors $\mathbf{r} = (r_1, ...r_k)$ with $r_i = (r_{i_t 1}, ...r_{i_t n})$, define the mean μ and covariance matrix Σ as

$$\mu = (\mu_1, ...\mu_k)$$

$$\Sigma = \text{Cov}(r_i, r_j)_{1 \le i \le j \le k}$$

Markowitz portfolio theory states that the efficient portfolio which gives the maximum risk-adjusted return for a given level of risk-aversion factor γ is the following:

$$\mathbf{w} = \text{argmax}_w w^T \mu_r - \gamma w^T \Sigma w$$

$$\text{subject to } w^T 1 = 1$$

In general, this quadratic optimization cannot be solved in closed form, and requires the use of a quadratic programmer. We used the *cvx* toolbox for MATLAB, and *quadprog* in R. The main goal of our work is to find plug-in estimates for the mean vector μ and covariance matrix Σ which vary over time.

B. Outline

The rest of the paper is organized as follows. In Section II, we describe the macroeconomic and sector return data being examined and demonstrate how Principle Components Analysis can reduce the dimension of the economic space. In Section III, we explore two economic proximity-based weighting scheme for estimating better mean vector and covariance matrix. The first approach identifies in discrete economic ”regimes” using k-means; we extended this idea to continuous space using a Gaussian kernel in the second approach. In Section IV, we depart from the Markowitz framework and reformulate portfolio construction as a reinforcement learning problem. Using fitted Q-iteration, we look for the optimal policy given a economic state. In Section V, we evaluate the performance of different strategies and compare to the benchmarks. Finally, Section VI concludes the paper.

II. Data Collection and PCA

A. Sector Indices

Since there are not enough sector/industry indices with long history (+10 years), we decide to construct a set of indices ourselves. Using Russell 1000 as an underlying universe, we group the index constituents by 2-digit Global Industry Classification Standard (GICS) for the period from January 1980 to March 2009. Then, we form 24 equal-weighted sector portfolios at the beginning of each month and calculate returns assuming 1-month holding period. Dividends are assumed to be reinvested to the stock on the ex-dividend date. The description of the sectors is shown in Appendix A.

In constructing the indices, we ensure that there is enough number of stocks in each sector so that index returns are not driven by a few stocks. Since Russell 1000 is a well-known large-cap index, we expect each constituent to have sufficient liquidity and low transaction cost for trade. In other words, we believe that these sector index returns can be replicated at reasonable cost in real practice.

B. Economic Data

The primary source of economic data is Federal Reserve Bank of St. Louis. The Fed collects numerous economic data from several sources and allows users to download them as text files. We gather a set of 27 economic indicators we believe has significant impact on certain sectors as well as on general economy. The list of macroeconomic information is shown in Appendix B.

In order to minimize look-ahead bias when trading, we lag these economic data by 1-2 months as appropriate. For example, we assume the GDP figure for 2009Q1 (which ended in March) became available for trading on May. Unfortunately, there is no easy solution for look-ahead bias due to revisions. As far as we know, there is no comprehensive point-in-time database that tracks the changes in macroeconomic indices.

C. Economic Dimension Reduction Using PCA

The economic time-series collected can all be thought of as proxies for the economic condition at time t, and although each is a different measure, it is reasonable to assume that they can be largely described using a lower dimensional space. To reduce complexity of optimizations in subsequent sections of this paper, and to gain an intuitive understanding of the underlying economic state space, we perform Principle Components Analysis (PCA) on the standardized indicator series to reduce the dimensionality.

Since our economic time-series data is of varying lengths, at each point in time for which we compute an economic decomposition, in order to use complete series', we select only a subset of our 27 economic indicators. This is achieved heuristically by optimizing the tradeoff being number of economic indicators, and length of the available series, by maximizing the product of these two quantities.

Having selected the economic indicators, we perform PCA, and choose the number of PCs that explain at least 80% of the variance in the economic condition as our model selection criteria. Using this decomposition, we can represent an economic state at time t as the projection of the observation at time t onto the first n PCs, we denote this state vector as $\mathbf{s_t}$.

Recall that our goal is describe the expected relationship between a group of sectors conditioned on the current economic condition. Towards this end, we would like for our economic decomposition in the reduced space to cluster when according to when economic condition is 'similar' in some sense. If we can get this, then we can reasonably expect that the mean μ_t and covariance Σ_t generated by weighting observations at time $s < t$ by the distance between the economic condition at time t to that at time s.

Towards this, we attempt to make sense of the PCA decomposition to determine if the results make sense fundamentally. Looking at the biplot relationship between the first two PCs in Figure 1, we find that the first PC, clearly indicates a clustering of projections of economic indicators in a sensible way. In the right half of the plot, the vectors are generally negative economic indicators (unrate, credit), while on the left half, the vectors are generally positive indicators (gdp, payrolls). The second PC seems to be related to an inflation metric, with high inflation represented on the bottom half, and low inflation on the top.

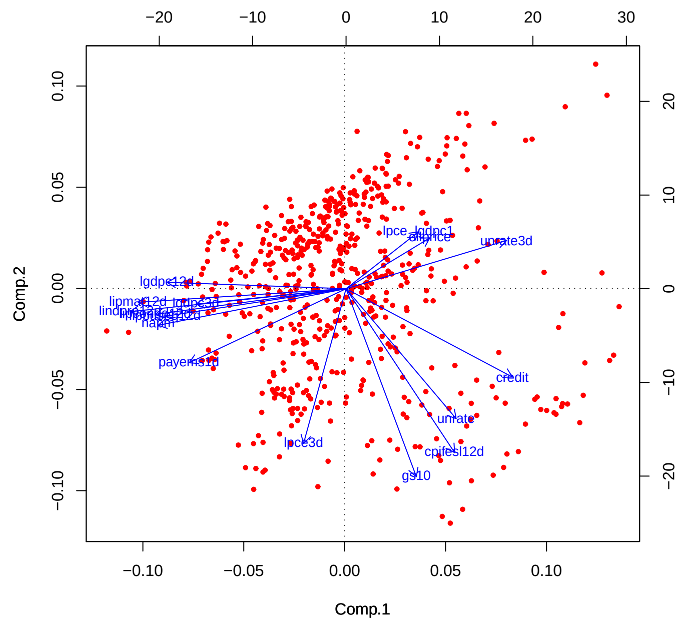

Fig. 1. Biplot relationship of 1st and 2nd PC

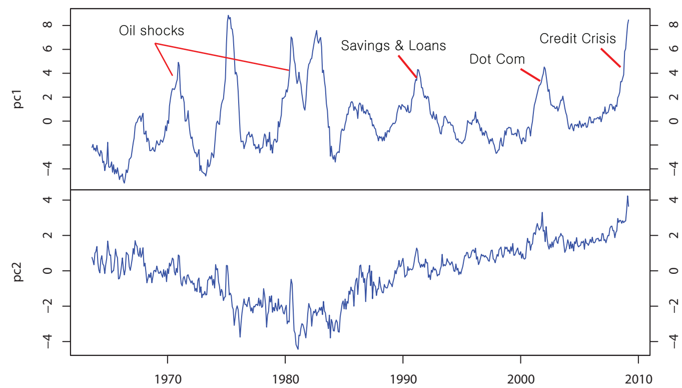

Fig. 2. Projection of Economic State onto First Two PC's

Figure 2 plots the projection of the economic state onto the first two principle components since the 1960's. Notice that the PC clearly identifies serious recessionary periods in history, from the Oil Shocks of the 1970's, to the Credit Crisis of 2008-2009. This validates that our economic decomposition is stable and is capable of identifying economic regimes. We now proceed to condition our portfolio optimization decisions on this information.

III. Economic Proximity Based Weighting

A. Discrete states: k-means clustering

In our first approach, motivated by prior work, we attempt to identify discrete economic regimes - e.g., *Good* vs. *Bad* economy in 2-state regime - within which we may expect the mean vector μ_t and covariance matrix Σ_t to be stationary. To do this, apply k-means clustering algorithm over all observations in the training window of m-months to cluster the economic condition in PCA space $(PC_1, PC_2, ..., PC_n)$ around k-centroids.

Next, we project the current economic state into PCA space to obtain $\mathbf{s_t}$, and attempt to determine if the current number of clusters is sufficient. We do this by testing the hypothesis that the current economic state is a new cluster, which is evaluated by comparing the silhouette statistic between the two models. If a new cluster is required, we create one, otherwise, we assign the current economic state to an existing cluster.

Finally, we compute the sample mean and covariance all observations that fall into the same cluster as the current economic state $\mathbf{s_t}$, and use the resulting mean vector μ_t and covariance matrix Σ_t in our convex optimization problem from Section 1. Please see the Evaluation section for results and discussion.

One observation made during this exploration was that it appears to take an arbitrarily large number of clusters to describe the n-dimensional PCA economic state space, as new observations frequently generated new clusters, which is causing us to lose a lot of training data. This motivates the idea of using all observations, but weighting them in continuous space.

B. Continuous State: Gaussian Kernel

A natural extension of the discrete regime idea in the previous section is, assuming that there are arbitarily many states in the economy, come up with a way to generate conditional mean vector μ_t and covariance matrix Σ_t using weighted sample statistics of all observations training window, weighted by some measure of how similar the current state $\mathbf{s_t}$ is to each of the previous $\mathbf{s_i}$.

This is achieved by weighting each observation by a Gaussian kernel, such that $w_i = exp(-(\mathbf{s_i} - \mathbf{s_t})^T \tau (\mathbf{s_i} - \mathbf{s_t}))$, where τ is a diagonal bandwidth matrix that tunes how quickly the weights decay. The value of τ was tuned by empirically maximizing the tradeoff between assigning near-zero weight to too many observations, and not clearly discriminating between economically dissimilar states. Figure 3 plots a sample of the weights for 1980 to 1991. Notice that this was a recessionary period, and higher weights are given to prior periods under which the market was in distress.
Next, since we want our portfolio to be fully invested, we normalize the weights w_i to sum to 1, and generate the conditional mean vector μ_t and covariance matrix Σ_t as

$$\mu_t = \frac{1}{m} \sum_{i=t-m}^{t} (r_i w_i)$$

$$\Sigma_t = \text{Cov}(r_{i,s} w_s, r_{j,s} w_s)_{t-m \le s \le t, 1 \le i \le j \le k}$$

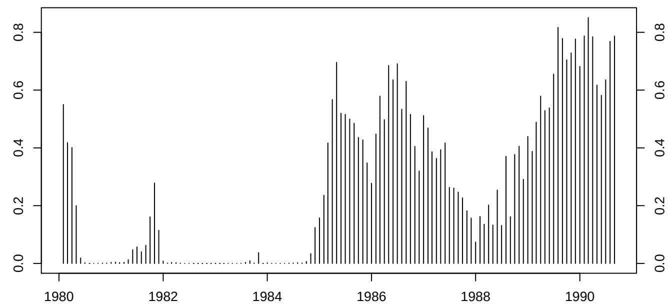

Fig. 3. Gaussian Kernel Weights w_i at t=1991

As before, we use estimates as plug-in estimates for the convex optimization problem posted at the start. Please see the Evaluation section for results and discussion.

IV. Reinforced Learning: Fitted Q-iteration

While the previous approaches have appealing economic logic and supportive result, they are still based on and, in a sense, confined by Markowitz framework. Departing from this classical theory, we reformulate portfolio construction decision as a reinforced learning problem whose objective is to find the optimal sector weighting policy given the current economic state. In doing so, the immediate difficulty we face is that both our state and action spaces are continuous, high-dimensional spaces; although reduced in dimension through PCA, economic state $\mathbf{S}$ still requires 3 5 dimensions whereas action space $\mathbf{A}$ lies in 24 dimensional constrained space. To solve this issue, we resort to a variant of fitted Q-iteration (FQI) called advantaged-weighted regression (FQI-AWR). [6]

A. FQI using Advantage-Weighted Regression

FQI is an iterative algorithm to approximate the state-action value function $Q : \mathbf{S} \times \mathbf{A} \to \mathbb{R}$ using historical state-action pairs $(\mathbf{s}, \mathbf{a})$. The estimated Q-functions are searched over the action space $\mathbf{A}$ to find the optimal value function $\mathbf{V}(\mathbf{s})$. The updated value functions are subsequently used for finding better approximation of Q and the process repeats. However, during the value update, the greedy operation $\mathbf{V}(\mathbf{s}) = max_{\mathbf{a}} Q(\mathbf{s}, \mathbf{a})$ often becomes prohibitively expensive for the continuous action space. FQI-AWR handles this situation by replacing the costly max operator with the weighted regession estimate $\hat{\mathbf{V}}(\mathbf{s})$ whose weights are derived by advantage $\mathbf{A} = Q(\mathbf{s}, \mathbf{a}) - \mathbf{V}(\mathbf{s})$. More specifically, AWR algorithm performs updates according to the following equations:

Q-function update

$$Q_{k+1}(\mathbf{s}, \mathbf{a}) = \mathbf{s_A} ({S_A}^T W S_A)^{-1} {S_A}^T W Q_k$$

V-function update

$$V_{k+1}(\mathbf{s}) = \mathbf{s} (S^T U S)^{-1} S^T U Q_{k+1}$$

Policy update (after V_k converges)

$$\mu(\mathbf{s}) = \mathbf{s}(S^T US)^{-1} S^T UA$$

where $\mathbf{s_A} = [s^T, a^T]^T, S_A = [\mathbf{s_{A(1)}}, \mathbf{s_{A(2)}}, ..., \mathbf{s_{A(T)}}]$ is the state-action matrix with the associated diagonal weight matrix $W_{ii} = w_\mathbf{s}(\mathbf{s_i})w_\mathbf{a}(\mathbf{a_i})$. $S = [\mathbf{s}_1^T, \mathbf{s}_2^T, ..., \mathbf{s}_T^T]$ is the state-only matrix whose weight is given by $U_{ii} = w(\mathbf{s})u(\mathbf{s}, \mathbf{a})$ where u is the normalized advantage. $A = [\mathbf{a}_1^T, \mathbf{a}_2^T, ..., \mathbf{a}_T^T]$ represents the action matrix. [6] Notice that, we need a distance metric for the action space in order to compute W. As in the economic state case, we use Gaussian kernel.

The AWR procedures described above have two very intuitive features that differentiate itself from the previous approaches. First, it measures the similarity based on both state and action. This enables us to appropriately evaluate the cases where economic situations are similar, but the actions taken are quite different. Secondly, the value updating step based on normalized advantage inherently awards the good actions (i.e., making positive returns) while penalizing the bad actions. For instance, suppose we have an observation in history that is economically similar to the current state, but the action taken at that time resulted in a very negative return. The natural response is not to repeat the same mistake and AWR updating policy exactly perform this task.

B. Implementation issues

FQI-AWR framework is straightforward and, for the most part, directly applicable to our setting. However, there are a few places where we make implicit assumptions or deviate from the standard model as we see necessary. We clarify these points here. First, we set our initial policy as the benchmark Markowitz weights plus a random turbulence term. The random turbulence is taken from a normal distribution $\mathcal{N}(0, 0.05^2)$. Adding turbulence allows to look at the paths different from Markowitz portfolios. Second, our reward function is solely the return generated from the action chosen. We do not consider the risk implied by Markowitz model. Third, during the final step of applying the optimal policy to the current state, we rescale the weight vector $\mathbf{s}(S^T US)^(-1) S^T U$ so that it sums up to 1. While it is unconventional, we need this additional step to ensure fully invested portfolios. Conceptually, this is equivalent to adjust the bandwidth of locally weighted regression.

V. Evaluation

We evaluate the efficacy of our economic similarity-based methods against two benchmark portfolios: The first is the equal-weighted index, which is equivalent to invest the same proportion of money to each stock in the universe. The second benchmark is the portfolio formed by Markowitz optimization with no economic prior. We simply take the past 60-months of return and calculate mean and covariance matrix with equal weights. Moreover, for any portfolio that goes through Markowitz optimization, we impose two constraints: (1) the weights add up to 1 and (2) the maximum bet of ±5% from the equal-weighted portfolio. The backtest result for these five portfolios is shown in Figure 4 and in Table I.

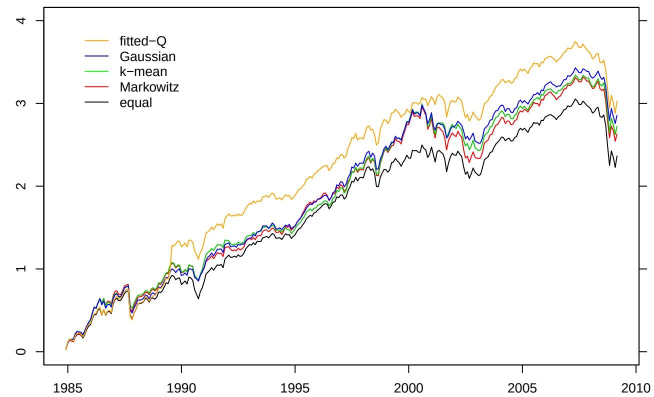

Fig. 4. Backtest Result from Jan 1985 to Mar 2009

	Equal	Markowitz	k-means	Gaussian	Fitted-Q
Ann. Return	10.2%	11.4%	11.8%	12.5%	13.3%
Info. Ratio	0.56	0.59	0.63	0.70	0.66

TABLE I
Backtest Performance Returns and Information Ratios

All of our three methods outperformed the benchmarks both in term of annualized return. We repeat backtests with varying combination of parameters, such as maximum bet, bandwidth for Gaussian kernel, etc. Although the result of FQI portfolio may vary stochastically due to the turbulence term, the overall performance of economic similarity-based methods is clearly better than the benchmarks.

VI. Conclusion

Our research is based on the presumption that sectors behaves in similar ways under similar economic conditions. To test our hypothesis, we first apply PCA to a set of macroeconomic variables to reduce the dimension of economic states. The resulting principle components successfully captures the major economic trends as well as other aspects of economy, e.g., inflation concern or credit-tightening.

We compare three methods of incorporating economic similarities into portfolio construction: (1) k-means clustering, (2) Gaussian kernel distance, and (3) fitted Q-iteration (FQI). While all of the three methods outperformed the benchmarks, the first two are a modification of classical Markowitz model. In the calculation of expected return and covariance matrix, in a departure form the assumption that these quantities are stationary, we compute weighted sample statistics, giving more weight to the observations which correspond to a similar economic state. The third approach poses the question into a reinforced learning framework, directly searching for the optimal policy using simulated strategy-return pairs. To cope with high dimensional, continuous state-action space, we adopt advantage-weighted regression technique that expresses the optimal policy as a weighted-sum of previous actions. Since we only reward return and do not penalize risk (covariance), the resulting policy gives the highest return but the information ratio less than the second method.

Given the significant improvement in our result, we believe that, to a certain degree, history does repeat and portfolio optimization can benefit from not only taking economic similarities into consideration but also evaluating the effectiveness of the actions taken under those circumstances.

VII. References

[1] N. Chen, R. Roll, and S. Ross, "Economic forces and the stock market," *Journal of Business*, vol. 59, 1986.

[2] J. Stangl, B. Jacobsen, and N. Visaltanachoti, "Sector rotation over business-cycles," Mar. 2008.

[3] R. Madrid, K. McNeil, J. Miller, E. Slover, and B. Zimmerman, "Economic forces and stock sectors," Dec. 2008.

[4] J. Shanken and M. I. Weinstein, "Economic forces and the stock market revisited," *Journal of Empirical Finance*, Sep. 2005.

[5] T. L. Lai and H. Xing, *Statistical Models and Methods for Financial Markets*. New York, NY: Springer, 2008.

[6] G. Neumann and J. Peters, "Fitted q-iteration by advantage weighted regression," *Advances in Neural Information Processing Systems 22 (NIPS 2008)*, 2009.

VIII. Appendices

A. Appendix A - Underlying Sectors

GICS4 Code	Description
1010	Energy
1510	Materials
2010	Capital Goods
2020	Commercial Services & Supplies
2030	Transportation
2510	Automobiles & Components
2520	Consumer Durables & Apparel
2530	Hotels, Restaurants & Leisure
2540	Media
2550	Retailing
3010	Food & Staples Retailing
3020	Food, Beverage & Tobacco
3030	Household & Personal Products
3510	Health Care Equipment & Services
3520	Pharmaceuticals & Biotechnology
4010	Banks
4020	Diversified Financials
4030	Insurance
4040	Real Estate
4510	Software & Services
4520	Technology Hardware & Equipment
4530	Semiconductors & Semiconductor Equipment
5010	Telecommunication Services
5510	Utilities

TABLE II
UNDERLYING SECTORS - SOURCE: REUTERS

B. Appendix B - Economic Indicators

Name	Description
umcsent	University of Michigan: Consumer Sentiment
umcsent3d	3-month Change of umcsent
lgdpc12d	12-month Change in log of GDP
lgdpc3d	3-month Change in log of GDP
tdsp	Household Debt Service Payments
tdsp3d	3-month Change of tdsp
lpce3d	3-month Change in Log of Personal Cons. Exp.
lpce_logdpc1	log(pce) - log(gdpc1)
unrate	Civilian Unemployment Rate
unrate3d	3-month Change of Unemployment Rate
payems1d	1-month Change of Nonfarm Payrolls
credit	Credit Spread (Baa - Aaa)
term	30-Year Treasury- 10-Year Treasury CM Rate
dfedtar	Federal Funds Target Rate
gs10	10-Year Treasury Constant Maturity Rate
oilprice	Oil Price - West Texas Intermediate
cpifesl12d	CPI All Items Ex Food & Energy - 12-month Change
dgorder12d	Durable Goods Order - 12-month Change
dgorder3d	Durable Goods Order - 3-month Change
napm	ISM Manufacturing: PMI Composite Index
tcu	Capacity Utilization: Total Industry
lindpro12d	Industrial Production Index
lipmat12d	Industrial Production: Materials
liputil12d	Industrial Production: Electric and Gas Utilities
lipmine12d	Industrial Production: Mining
lipbuseq12d	Industrial Production: Business Equipment
lipcongd12d	Industrial Production: Consumer Goods

TABLE III
ECONOMIC INDICATORS - SOURCE: ST. LOUIS FED