

Detecting Changes in Asset Co-Movement Using the Autoencoder Reconstruction Ratio

Bryan Lim, Stefan Zohren, and Stephen Roberts

Abstract—Detecting changes in asset co-movements is of much importance to financial practitioners, with numerous risk management benefits arising from the timely detection of breakdowns in historical correlations. In this article, we propose a real-time indicator to detect temporary increases in asset co-movements, the Autoencoder Reconstruction Ratio, which measures how well a basket of asset returns can be modelled using a lower-dimensional set of latent variables. The ARR uses a deep sparse denoising autoencoder to perform the dimensionality reduction on the returns vector, which replaces the PCA approach of the standard Absorption Ratio (Kritzman *et al.*, 2011), and provides a better model for non-Gaussian returns. Through a systemic risk application on forecasting on the CRSP US Total Market Index, we show that lower ARR values coincide with higher volatility and larger drawdowns, indicating that increased asset co-movement does correspond with periods of market weakness. We also demonstrate that short-term (i.e. 5-min and 1-hour) predictors for realised volatility and market crashes can be improved by including additional ARR inputs.

INTRODUCTION

The time-varying nature of asset correlations has long been of much interest to financial practitioners, with short-term increases in the co-movement of financial returns widely documented around periods of market stress (Preis *et al.*, 2012; Packham & Woebeking, 2019; Campbell *et al.*, 2002; Cappiello *et al.*, 2006; Billio *et al.*, 2010). From a portfolio construction perspective, short-term increases in asset correlations have frequently been linked to spikes in market volatility (Campbell *et al.*, 2002) – with explanatory factors ranging from impactful news announcements (At-Sahalia & Xiu, 2016) to “risk-on risk-off” effects (Dungey *et al.*, 2018) – indicating that diversification can breakdown precisely when it is needed the most. In terms of systemic risk, increased co-movement between market returns is viewed as a sign of market fragility (Kritzman *et al.*, 2011), as the tight coupling between markets can deepen drawdowns

when they occur, resulting in financial contagion. The development of methods to detect real-time changes in returns co-movement can thus lead to improvements in multiple risk management applications.

Common approaches for co-movement detection can be broadly divided into two categories. Firstly, multivariate Gaussian models have been proposed to directly model time-varying correlation dynamics – with temporal effects either modelled explicitly (At-Sahalia & Xiu, 2016; Dungey *et al.*, 2018), or incorporated simply by recalibrating covariance matrices daily using a sliding window of data (Preis *et al.*, 2012). More adaptive correlation estimates can also be provided by calibrating models using high-frequency data – as seen in the multivariate HEAVY models of Noureldin *et al.* (2012). However, several challenges remain with the modelling approach. For instance, naïve model specification can lead to spurious statistical effects (Loretan & English, 2000) reflecting correlation increases even when underlying model parameters remain unchanged. Furthermore, given that correlations are modelled in a pair-wise fashion between assets, condensing a high-dimensional correlation matrix into a single number measuring co-movement changes can be challenging in practice.

An alternative approach adopts the use of statistical factors models to capture common movements across the entire portfolio, typically projecting returns on low-dimensional set of latent factors using principal component analysis (PCA) (Kritzman *et al.*, 2011; Billio *et al.*, 2010; Zheng *et al.*, 2012). The popular Absorption Ratio (AR), for example, is defined by the fraction of the total variance of assets explained or absorbed by a finite set of eigenvectors (Kritzman *et al.*, 2011) – which bears similarity to metrics to evaluate eigenvalue significance in other domains (e.g. the Fractional Spectral Radius of Rezek & Roberts (1998)). Despite their ability to quantify co-movement changes with a single metric, with a larger AR corresponding to increased co-movement, PCA-based indicators require a covariance matrix to be estimated over a historical lookback window. This approach can be data-intensive for applications with many assets, as a long estimation window is required to ensure non-singularity

B. Lim, S. Zohren and S. Roberts are with the Department of Engineering Science and the Oxford-Man Institute of Quantitative Finance, University of Oxford, Oxford, United Kingdom (email: {blim, zohren, sjrob}@robots.ox.ac.uk).

of the covariance matrix (Billio *et al.*, 2010). Moreover, non-Gaussian financial returns (Jondeau *et al.*, 2007) can violate the normality assumptions required by PCA, potentially making classical PCA unsuitable for high-frequency data. While alternatives such as Independent Component Analysis (ICA) have been considered for modelling non-Gaussian financial time series data (Shah & Roberts, 2013), they typically maintain the assumption that assets are linear combinations of driving factors – and potential improvements can be made using non-linear approaches. As such, more adaptive non-parametric methods are hence needed to develop real-time indicators for changes in co-movement.

Advances in deep learning have demonstrated the benefits of autoencoders architectures for dimensionality reduction in complex datasets (Goodfellow *et al.*, 2016). In particular, autoencoders have had notable successes in feature extraction in images (Cho, 2013; Liu & Zhang, 2018; Freiman *et al.*, 2019), vastly out-performing traditional methods for dimensionality reduction in non-linear datasets. More recently, autoencoders have been explored as a replacement for PCA in various financial applications, encoding latent factors which account for non-linearities in return dynamics and allow for conditioning on exogenous covariates. Gu *et al.* (2019), for instance, use a series of encoders to estimate both latent factors and factor loadings (i.e. betas) used in conditional asset pricing models, demonstrating better out-of-sample pricing performance compared to traditional linear methods. Kondratyev (2018) explores the use of autoencoders to capture a low-dimensional representation of the term structure of commodity futures, and also shows superior reconstruction performance versus PCA. Given the strong performance enhancements of autoencoder architectures on non-linear datasets, replacing traditional factor models with autoencoders could be promising for co-movement measurement applications as well.

In this article, we introduce the Autoencoder Reconstruction Ratio (ARR) – a novel indicator to detect co-movement changes in real-time – based on the average reconstruction error obtained from applying autoencoders to high-frequency returns data. Adopting the factor modelling approach, the ARR uses deep sparse denoising autoencoders to project second-level intraday returns onto a lower-dimension set of latent variables, and aggregates reconstruction errors up to the desired frequency (e.g. daily). Increased co-movement hence corresponds to periods where reconstruction error is low, i.e. when returns are largely accounted for by the autoencoder’s latent factors. In line with the canonical Absorption Ratio,

we evaluate the performance of the ARR by considering an application in systemic risk – using increased co-movement of sector index returns to improve volatility and drawdown predictions for the total market. The information content of the ARR is evaluated by measuring the performance improvements of machine learning predictors when the ARR is included as a covariate. Based on experiments performed for 4 prediction horizons (i.e. 5-min, 1-hour, 1-day, 1-week), the ARR was observed to significantly improve performance for volatility and market crashes forecasts over short-term (5-min and 1-hour) horizons.

PROBLEM DEFINITION

For a portfolio of N assets, let $r(t, n)$ be the returns of the n -th asset at time t defined as below:

$$r(t, n) = \log p(t, n) - \log p(t - \Delta t, n), \quad (1)$$

where $p(t, n)$ is the value or price of asset n at time t , and Δt is a discrete interval corresponding to the desired sampling frequency.

In their most general form, statistical factor models map a common set of K latent variables to the returns of each asset as below:

$$r(t, n) = \bar{r}(t, n) + \epsilon(t, n) \quad (2)$$

$$= f_n(\mathbf{Z}(t)) + \epsilon(t, n), \quad (3)$$

where the residual $\epsilon(t, n) \sim N(0, \xi_n^2)$ is the idiosyncratic risk of a given asset, and $f_n(\cdot)$ is a generic function, and $\mathbf{Z}(t) = [z(t, 1), \dots, z(t, K)]^T$ is a vector of common latent factors.

Traditional factor models typically adopt a simple linear form, using asset-specific coefficients as below:

$$f_n(\mathbf{Z}(t)) = \beta(n)^T \mathbf{Z}(t), \quad (4)$$

where $\beta(n) = [\beta(n, 1), \dots, \beta(n, K)]^T$ is a vector of factor loadings. While various approaches are available for latent variable estimation – such as independent component analysis (ICA) (Fabozzi *et al.*, 2015) or PCA (Billio *et al.*, 2010) – the number of latent variables are typically kept low to reduce the dimensionality of the dataset (i.e. $K \ll N$).

Absorption Ratio

A popular approach to measuring asset co-movement is the Absorption Ratio (Kritzman *et al.*, 2011), which performs dimensionality reduction on returns using PCA. Co-movement changes are then measured based on the total variance absorbed by a finite set of eigenvectors.

This number K set in the original paper to be a fifth of the number N of available assets:

$$\text{AR}(t) = \frac{\sum_{k=1}^K \sigma_{E_k}^2}{\sum_{n=1}^N \sigma_{A_n}^2}, \quad (5)$$

where $\text{AR}(t)$ is the Absorption Ratio at time t , $\sigma_{E_k}^2$ is the variance of the k -th eigenvector, and $\sigma_{A_n}^2$ is the variance of the n -th asset.

AUTOENCODER RECONSTRUCTION RATIO

To detect changes in asset co-movements, we propose the Autoencoder Reconstruction Ratio below, which can be interpreted as the normalised reconstruction mean squared error (MSE) of high-frequency returns within a given time interval:

$$\text{ARR}(t, \Delta t) = \frac{\sum_{\tau=t-\Delta t}^t \sum_{n=1}^N (r(\tau, n) - \bar{r}(\tau, n))^2}{\sum_{\tau=t-\Delta t}^t \sum_{n=1}^N r(\tau, n)^2}, \quad (6)$$

where Δt is the time interval matching the desired sampling frequency, and $\bar{r}(\tau, n)$ is the return of asset n reconstructed by a deep sparse denoising autoencoder (see sections below). For our experiments, we train the autoencoder using one-second returns and compute ARR's across four different sampling frequencies (i.e. 5-min, 1-hour, 1-day, 1-week).

Relationship to the Absorption Ratio

The ARR can also be interpreted as a slight reformulation of the standard Absorption Ratio, and we examine the relationship between the two in this section.

Reintroducing the linear Gaussian assumptions used by the Absorption Ratio, we note that the denominator of Equation (6) contains a simple estimator for the realised variance (RV) of each asset (Barndorff-Nielsen & Shephard, 2002), i.e.:

$$\text{RV}(t, \Delta t, n) = \sum_{\tau=t-\Delta t}^t r(\tau, n)^2 \quad (7)$$

$$\approx \sigma_{A_n}^2. \quad (8)$$

Moreover, by comparing to Equation (2), we can see that the numerator can be interpreted as an estimate of the sum of residual variances, leading to the form of the ARR below:

$$\text{ARR}(t, \Delta t) = \frac{\sum_{n=1}^N \xi_n^2}{\sum_{n=1}^N \sigma_{A_n}^2}. \quad (9)$$

Assuming that residuals are uncorrelated between assets, we note that the sum of residual variances essentially

corresponds to the variance unexplained by the selected factors. The ARR can in this case be expressed as:

$$\text{ARR}(t, \Delta t) = \frac{\sum_{n=1}^N \sigma_{A_n}^2 - \sum_{k=1}^K \sigma_{E_k}^2}{\sum_{n=1}^N \sigma_{A_n}^2} \quad (10)$$

$$= 1 - \text{AR}(t) \quad (11)$$

Autoencoder Architecture

We adopt a deep sparse denoising autoencoder architecture (Goodfellow *et al.*, 2016) to compute reconstruction errors for the ARR. The network is predominantly divided into two parts – 1) a decoder which reconstructs the returns vector from a reduced set of latent factors, and 2) an encoder which performs the low-dimensional projection – both of which are described below.

Decoder:

$$\bar{\mathbf{r}}(t) = \mathbf{W}_1 \mathbf{h}_1(t) + \mathbf{b}_1 \quad (12)$$

$$\mathbf{h}_1(t) = \text{ELU}(\mathbf{W}_2 \mathbf{Z}(t) + \mathbf{b}_2) \quad (13)$$

where $\bar{\mathbf{r}}(t) = [\bar{r}(t, 1), \dots, \bar{r}(t, N)]^T$ is the vector of reconstructed returns, $\text{ELU}(\cdot)$ is the exponential linear unit activation function (Clevert *et al.*, 2016), $\mathbf{h}_1(t) \in \mathbb{R}^H$ is the hidden state of the decoder network, and $\mathbf{W}_1 \in \mathbb{R}^{N \times H}$, $\mathbf{W}_2 \in \mathbb{R}^{H \times K}$, $\mathbf{b}_1 \in \mathbb{R}^N$, $\mathbf{b}_2 \in \mathbb{R}^H$ are its weights and biases.

Encoder:

$$\mathbf{Z}(t) = \text{ELU}(\mathbf{W}_3 \mathbf{h}_2(t) + \mathbf{b}_3) \quad (14)$$

$$\mathbf{h}_2(t) = \text{ELU}(\mathbf{W}_4 \mathbf{x}(t) + \mathbf{b}_4) \quad (15)$$

where $\mathbf{Z}(t) \in \mathbb{R}^K$ is a low-dimensional projection of inputs $\mathbf{x}(t) = [r(t, 1), \dots, r(t, N), \mathcal{T}]^T$, $\mathbf{h}_2(t) \in \mathbb{R}^H$ is the hidden state of the encoder network, and $\mathbf{W}_3 \in \mathbb{R}^{K \times H}$, $\mathbf{W}_4 \in \mathbb{R}^{H \times N}$, $\mathbf{b}_3 \in \mathbb{R}^K$, $\mathbf{b}_4 \in \mathbb{R}^H$ are its weights and biases. We note that the time-of-day \mathcal{T} – recorded as the number of seconds from midnight – is also included in the input along with sector returns, allowing the autoencoder to account for any intraday seasonality present in the dataset.

To ensure that dimensionality is gradually reduced, we set $K = \lfloor N/5 \rfloor$ as per the original Absorption Ratio paper, and fix $H = \lfloor (N + K)/2 \rfloor$.

Network Training

To improve generalisation on test data, sparse denoising autoencoders introduce varying degrees of regularisation.

Firstly, a sparsity penalty in the form of a L1 regularisation term is added to the reconstruction loss as below:

$$\mathcal{L}(\Theta) = \frac{\sum_{t=1}^T \sum_{n=1}^N (r(t, n) - r(\bar{t}, n))^2}{TN} + \alpha \|Z_t\|_1, \quad (16)$$

where Θ represents all network weights, α is a penalty weight which we treat as a hyperparameter, and $\|\cdot\|_1$ is the L1 norm. Secondly, inputs are corrupted with noise during training – forcing the autoencoder to learn more general relationships from the data. Specifically, we adopt masking noise for our network, which simply corresponds to the application of dropout (Srivastava *et al.*, 2014) to encoder inputs.

Hyperparameter optimisation is performed using 20 iterations of random search, with networks trained up to a maximum of 100 epochs per search iteration. Full hyperparameter details are listed in the appendix for reference.

FORECASTING SYSTEMIC RISK WITH THE ARR

To demonstrate the utility of the ARR, we consider applications in systemic risk forecasting – using ARRs computed from subsector indices to predict risk metrics associated with the overall market. Specifically, we consider the two key use-cases for the ARR:

- 1) **Volatility Forecasting** – i.e. predicting market turbulence as measured by spikes in realised volatility;
- 2) **Predicting Market Crashes** – i.e. providing an early warning signal for sudden market declines, allowing for timely risk management.

Description of Dataset

We focus on the total US equity market and 11 constituent sector indices for our investigation, using high frequency total returns sampled every second to compute the ARR. Intraday index data from the Center of Research in Security Prices (CRSP) was downloaded via Wharton Research Data Services (WRDS (2019)) from 2012-12-07 to 2019-03-29, with the full list of indices provided in Exhibit 1.

Reconstruction Performance

We compare the reconstruction accuracy of the deep sparse denoising autoencoder and standard PCA to evaluate the benefits of a non-linear approach to dimensionality reduction. To train the autoencoder, we divide the data into a training set used for network backpropagation (2012-2014), a validation set used for hyperparameter optimisation (2015 only), and a out-of-sample test set

Ticker	Index Description
CRSPTMT	CRSP US Total Market Total-Return Index
CRSPRET	CRSP US REIT Total-Return Index
CRSPENT	CRSP US Oil and Gas Total-Return Index
CRSPMTT	CRSP US Materials Total-Return Index
CRSPIDT	CRSP US Industrials Total-Return Index
CRSPCGT	CRSP US Consumer Goods Total-Return Index
CRSPHCT	CRSP US Health Care Total-Return Index
CRSPCST	CRSP US Consumer Services Total-Return Index
CRSPTET	CRSP US Telecom Total-Return Index
CRSPUTT	CRSP US Utilities Total-Return Index
CRSPFNT	CRSP US Financials Total-Return Index
CRSPITT	CRSP US Technology Total-Return Index

Exhibit 1: US Market & Sector Indices Used in Tests

(2016-2019) used to quantify the information content of the ARR. PCA factors were calibrated using covariance matrix estimated with data from 2012-2016, keeping the same out-of-sample data as the autoencoder. The returns vector is projected onto 2 latent variables, with the dimensionality of the latent space taken to be 1/5 the number of indices used – as per the original Absorption Ratio paper (Kritzman *et al.*, 2011).

We evaluate the reconstruction accuracy using the test dataset, based on the R-squared (R^2) of reconstructed returns for each prediction model. Given that the ARR is computed based on sector data alone, we only focus on the reconstruction accuracy of the 11 sector indices. We also test for the significance of the results with a bootstrap hypothesis test – using 500 bootstrap samples and adopting the null hypothesis that the autoencoder has a lower reconstruction R^2 than PCA.

	PCA	Autoencoder	P-Value
R^2	0.340	0.461*	< 0.01

Exhibit 2: Out-of-Sample Reconstruction R^2 (2016-2019)

From the results in Exhibit 2, we can see that the autoencoder greatly enhances reconstruction accuracy – increasing the out-of-sample R^2 by more than 35% and statistically significant with 99% confidence. These improvements highlight the benefits of adopting a non-linear approach to dimensionality reduction, and the suitability of the autoencoder for modelling intraday returns. Such an approach can have wider applications in various areas of quantitative finance.

Empirical Analysis

Next, we perform an exploratory investigation into the relationships between the total market index, and the

ARR of its constituent sectors. Specifically, we examine the following metrics aggregated over different sampling frequencies ($\Delta t \in \{5\text{-min}, 1\text{-hour}, 1\text{-day}, 1\text{-week}\}$):

- **Returns** – Corresponding to the difference in log prices based on Equation (1).
- **Log Realised Volatility (Log RV)** – Computed based on the logarithm of the simple realised volatility estimator of Equation (7).
- **Drawdowns (DD)** – Determined by the fractional decrease from the maximum index value from the start of our evaluation period.

ARRs were similarly computed for the 4 sampling frequencies, and aggregated based on Equation (6). Using the KDE plots of Exhibits 3 to 5, we perform an in-sample (2012-2015) empirical analysis of the coincident relationships between the ARR and the metrics above. To avoid spurious effect, we remove the eves of public holidays from our analysis – where markets are open only in the morning and trading volumes are abnormally low. In addition, we winsorise the data at the 1st and 99th percentiles, reducing the impact of outliers to improve data visualisation.

Looking at the KDE plots for returns against the ARR in Exhibit 3, we observe a noticeable increase in the dispersion of returns appears at low ARR values. This is echoed by the KDE plots for Log RV in 4 – which appears to increase linearly with decreasing ARR. Drawdowns behave in a similar fashion as well, with larger drawdowns appearing to occur at low values of the ARR. Effects also are consistent across different horizons, with similar patterns observed for all sampling frequencies. On the whole, the results validate the findings observed in previous works – indicating that increased asset co-movements do indeed coincide with periods of market weakness – with lower ARR values observed around periods of high volatility or high drawdowns.

To visualise how the ARR changes over time, we plot 5-min ARR (i.e. $ARR(t, 5\text{-min})$) against the prices and drawdowns of the CRSP Total Market Index in Exhibit 6. Given the noisiness of the raw ARR values, we also included a smoothed version of 5-min ARR in the bottom – computed based on an exponentially weighted moving average with a 1-day half-life. We can see from the results that dips in the ARR occur slightly before large drawdown periods, particularly around the sell-offs of August 2015 and early 2018. As such, the ARR could potentially be used to improve predictions of volatility spikes or sudden market crashes in the near future – which we further investigate in the next section.

QUANTIFYING THE INFORMATION CONTENT OF THE ARR

We attempt to quantify the information content of the ARR by observing how much it improves forecasting models for the risk metrics above. To do so, we adopt several machine learning benchmarks for risk prediction, and evaluate the models' forecasting performance both with and without the ARR included in its inputs. This allows us to evaluate how much additional information is provided by the ARR, above that provided by temporal evolution of realised volatility or drawdowns alone.

Benchmark Models

Given the non-linear relationships observed between returns and drawdowns versus the ARR, we adopt a series of machine learning benchmarks on top of linear models. Specifically, we consider 1) linear/logistic regression, 2) gradient boosted decision trees (GBDTs), and 3) simple multi-layered perceptrons (MLPs) for our forecasting applications. Hyperparameter optimisation is performed using up to 200 iterations of random search, with optimal hyperparameters selected using 3-fold cross validation. Additional training details can also be found in the appendix.

Volatility Forecasting Methodology

We treat volatility forecasting as a regression problem, focusing on predicting Log RV over different horizons – using a variant of the HAR-RV model of Corsi (2009). For 5-min Log RV, a simple HAR-RV model can take the form below, incorporating log RV terms for longer horizons into the forecast:

$$\begin{aligned} \bar{\nu}(t + 5\text{-min}, 5\text{-min}) &= \beta_1 \nu(t, 5\text{-min}) + \beta_2 \nu(t, 1\text{-hour}) \\ &+ \beta_3 \nu(t, 1\text{-day}) + \beta_4 \nu(t, 1\text{-week}), \end{aligned} \quad (17)$$

where $\bar{\nu}(t + 5\text{-min}, 5\text{-min})$ is 5-min log RV predicted over the next time-step, $\nu(t, \Delta t)$ is the log RV from $t - \Delta t$ to t , and $\beta_{(\cdot)}$ are linear coefficients.

We adopt a similar non-linear variant for our forecast for each prediction horizon δ . For tests with the ARR, this takes the form:

$$\bar{\nu}(t + \delta, \delta) = g_1(\psi(t), \omega(t)), \quad (18)$$

where $\delta \in \Psi$, $\Psi = \{5\text{-min}, 1\text{-hour}, 1\text{-day}, 1\text{-week}\}$, $\psi(t) = \{\nu(t, d) : d \in \Psi \text{ and } d \geq \delta\}$, $\omega(t) = \{ARR(t, d) : d \in \Psi \text{ and } d \geq \delta\}$, and $g_2(\cdot)$ is a prediction model mapping inputs to returns. The combination of both realised volatility and ARR values hence allows

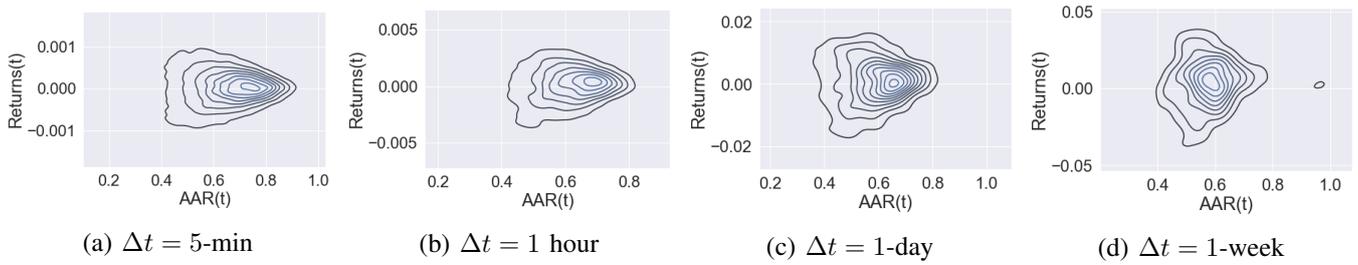


Exhibit 3: In-sample KDE Plots for Returns vs. ARR.

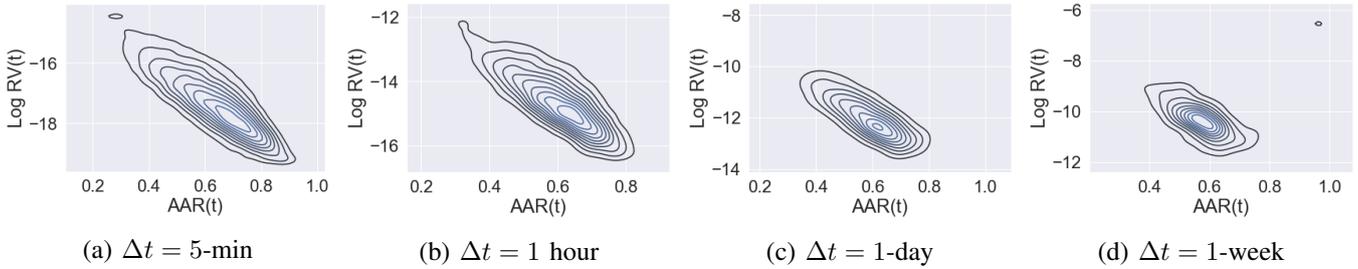


Exhibit 4: In-Sample KDE Plots for Log RV vs. ARR.

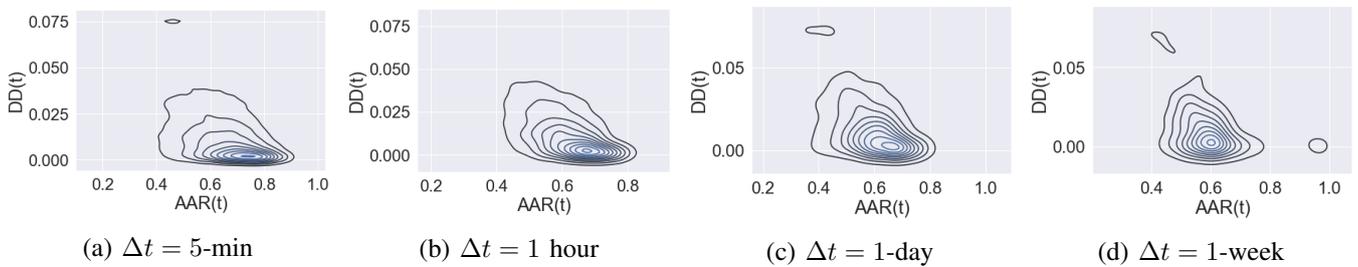


Exhibit 5: In-Sample KDE Plots for Drawdowns vs. ARR.

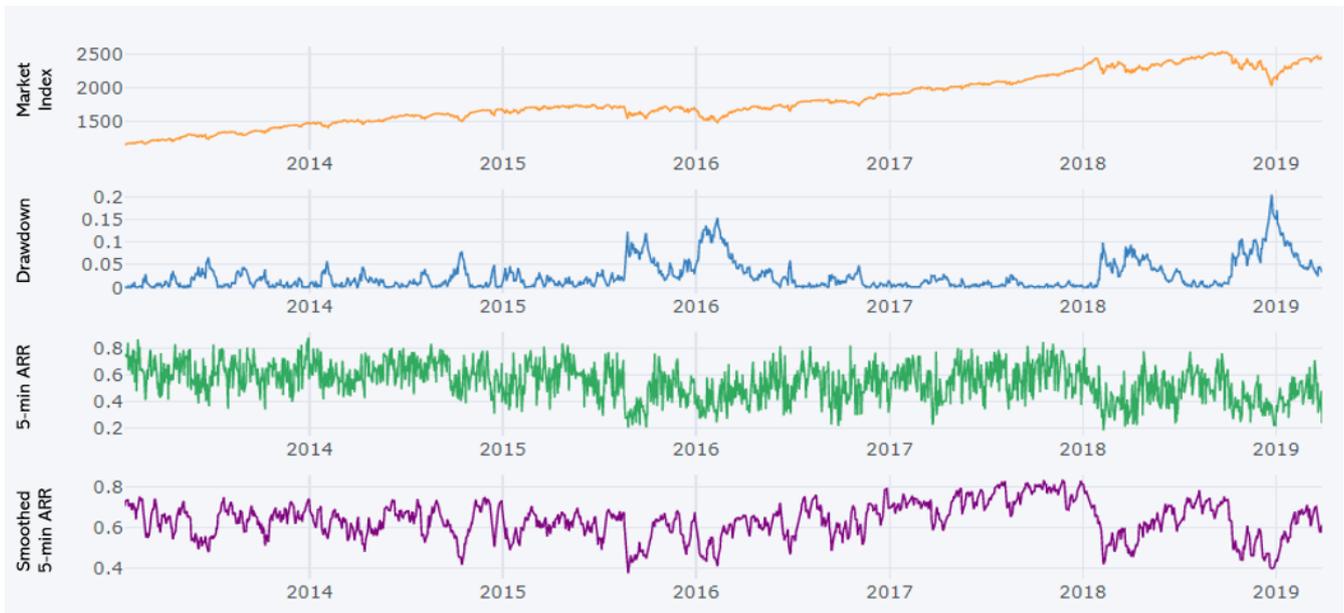


Exhibit 6: 5-min ARR from 2012 to 2019.

us to determine if the ARR supplies any information above that provide the volatility time series alone.

For tests without the ARR, we continue to use Equation (18) but omit ARR values, i.e. $\omega(t) = \emptyset$. We evaluate regression performance using the R^2 of forecasted log RV.

Crash Prediction Methodology

We take a binary classification approach to forecasting market crashes, using our benchmark models to predict the onset of a sharp drawdown. First, we define a z -score metric for returns $\gamma(t, \lambda)$ as below:

$$\gamma(t, \lambda) = \frac{r(t) - m(t, \lambda)}{s(t, \lambda)}, \quad (19)$$

where $r(t)$ is the return for the total market index, $m(t, \lambda)$ is the exponentially weighted moving average for returns with a half-life of λ , and $s(t, \lambda)$ its exponentially weighted moving standard deviation.

Next, based on our z -score metric, we define a market crash to be a sharp decline in market returns, i.e.:

$$c(t) = \mathbb{I}(\gamma(t, \lambda) < C), \quad (20)$$

where $c(t) \in \{0, 1\}$ is a crash indicator, $\mathbb{I}(\cdot)$ is an indicator function, and C is the z -score threshold for returns. For our experiments, we set λ to be 10 discrete time steps and $C = -1.5$.

We then model crash probabilities using a similar form to Equation (18):

$$p(c(t) = 1) = g_2(\psi(t), \omega(t)), \quad (21)$$

where $g_2(\cdot)$ is a function mapping inputs to crash probabilities $p(c(t) = 1)$.

Given that crashes are rare by definition – with $c(t) = 1$ for less than 10% of time steps for daily frequencies – we also oversample the minority class to address the class imbalance problem. Classification performance is evaluated using the area under the receiver operating characteristic (AUROC).

Results and Discussion

The results for both volatility forecasting and market crash prediction can be found in Exhibit 7, both including and excluding ARR values. To determine the statistical significance of improvements, we conduct a bootstrap hypothesis test under the null hypothesis that performance results are better when the ARR is included, using a non-parametric bootstrap with 500 samples.

From the volatility forecasting R^2 values in Exhibit 7a, we can see that the ARR consistently improves 5-min and 1-hour forecasts for all non-linear models, and

significant improvements are observed for linear models for 5-min sampling frequencies. This indicates that the ARR is informative for short-term forecasts, and can help enhance risk predictions in the near-term. For longer horizons however (i.e. 1-day and 1-week) we observe that the inclusion of the ARR reduces prediction accuracy – potentially indicating the presence of overfitting on the training set when ARRs are introduced.

For crash predictions AUROC results in Exhibit 7b, we note that the ARRs are observed to improve forecasts for all models and sampling frequencies – with statistical significance at the 99% level observed for both linear and GBDT forecasts over shorter horizons. This echoes the volatility forecasting results – indicating that ARRs can be useful to inform short-term risk predictions.

CONCLUSIONS

We introduce the Autoencoder Reconstruction Ratio (ARR) in this paper, using it as a real-time measure of asset co-movement. The ARR is based on the normalised reconstruction error of a deep sparse denoising autoencoder applied to a basket of asset returns – which condenses the returns vector onto a lower dimensional set of latent variables. This replaces the PCA modelling approach used by the Absorption Ratio of Kritzman *et al.* (2011), which allows the ARR to better model returns that violate basic PCA assumptions (e.g. non-Gaussian returns). Through experiments on a basket of 11 CRSP US sector indices, we demonstrate that the autoencoder significantly improves the out-of-sample reconstruction performance when compared to PCA, lowering the combined R^2 by more than 35%.

Given the links identified between increased asset co-movements and the fragility of the overall market in previous works (Kritzman *et al.*, 2011; Campbell *et al.*, 2002), we also evaluate the use of the ARR in a systemic risk application. First, we conduct an empirical analysis of the relationship between risk metrics of the combined market (using the CRSP US Total Market Index as a proxy) and the ARR computed from its sub sectors of the market. Based on an analysis of the KDE plots of risk metrics vs. ARRs, we show that low values of the ARR coincide with high volatility and high drawdown periods in line with previous findings. Next, we evaluate the information content of the ARR by testing how much the ARR improves risk predictions for a various benchmark models. We find that the ARR is informative for both volatility and market crash predictions over short horizons, and significantly increases 5-min and 1-hour forecasting performance across most model benchmarks.

		5-min	1-hour	1-day	1-week
Linear	With ARR	0.635*	0.496	0.389	0.579
	No ARR	0.626	0.536*	0.431*	0.611*
	P-Values	<0.01	>0.99	>0.99	>0.99
GBDT	With ARR	0.635*	0.554*	0.425	0.399
	No ARR	0.627	0.536	0.434	0.593*
	P-Values	<0.01	<0.01	0.821	>0.99
MLP	With ARR	0.641*	0.571*	0.387	0.527
	No ARR	0.631	0.549	0.426*	0.636*
	P-Values	<0.01	<0.01	>0.99	>0.99

(a) R^2 for Log RV Predictions

		5-min	1-hour	1-day	1-week
Linear	With ARR	0.598*	0.629*	0.585	0.372
	No ARR	0.587	0.570	0.582	0.333
	P-Values	<0.01	<0.01	0.464	0.284
GBDT	With ARR	0.590*	0.562*	0.529	0.605
	No ARR	0.567	0.516	0.465	0.575
	P-Values	<0.01	0.01	0.138	0.359
MLP	With ARR	0.589	0.564	0.586	0.423
	No ARR	0.588	0.558	0.540	0.210
	P-Values	0.374	0.392	0.234	<0.01

(b) AUROC for Crash Predictions

Exhibit 7: Out-of-Sample Forecasting Results With and Without ARR Inputs.

REFERENCES

- Abadi, Martín, et al. 2015. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org.
- At-Sahalia, Yacine, & Xiu, Dacheng. 2016. "Increased correlation among asset classes: Are volatility or jumps to blame, or both?," *Journal of Econometrics*, **194**(2), 205–219.
- Barndorff-Nielsen, Ole E., & Shephard, Neil. 2002. "Estimating quadratic variation using realized variance," *Journal of Applied Econometrics*, **17**(5), 457–477.
- Billio, Monica, Getmansky, Mila, Lo, Andrew W., & Pelizzon, Lorian. 2010. "Measuring systemic risk in the finance and insurance sectors," *MIT Sloan Research Paper No. 4774-10*. <http://ssrn.com/abstract=1571277>.
- Campbell, Rachel, Koedijk, Kees, & Kofman, Paul. 2002. "Increased correlation in bear markets," *Financial Analysts Journal*, **58**(1), 87–94.
- Cappiello, Lorenzo, Engle, Robert F., & Sheppard, Kevin. 2006. "Asymmetric Dynamics in the Correlations of Global Equity and Bond Returns," *Journal of Financial Econometrics*, **4**(4), 537–572.
- Cho, Kyung Hyun. 2013. "Simple sparsification improves sparse denoising autoencoders in denoising highly noisy images," *In: Proceedings of the 30th International Conference on Machine Learning*. ICML 2013.
- Clevert, Djork-Arne, Unterthiner, Thomas, & Hochreiter, Sepp. 2016. "Fast and accurate deep network learning by exponential linear units (ELUs)," *In: International Conference on Learning Representations*. ICLR 2016.
- Corsi, Fulvio. 2009. "A Simple Approximate Long-Memory Model of Realized Volatility," *Journal of Financial Econometrics*, **7**(2), 174–196.
- Dungey, Mardi, Erdemlioglu, Deniz, Matei, Marius, & Yang, Xiye. 2018. "Testing for mutually exciting jumps and financial flights in high frequency data," *Journal of Econometrics*, **202**(1), 18 – 44.
- Fabozzi, Frank, Giacometti, Rosella, & Tsuchida, Naoshi. 2015. *The ICA-based Factor Decomposition of the Eurozone Sovereign CDS Spreads*. IMES Discussion Paper Series 15-E-04. Institute for Monetary and Economic Studies, Bank of Japan.
- Freiman, Moti, Manjeshwar, Ravindra, & Goshen, Liran. 2019. "Unsupervised abnormality detection through mixed structure regularization (MSR) in deep sparse autoencoders," *Medical Physics*, **46**(5), 2223–2231.
- Goodfellow, Ian, Bengio, Yoshua, & Courville, Aaron. 2016. "Autoencoders," *Chap. 14 of: Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- Gu, Shihao, Kelly, Bryan T., & Xiu, Dacheng. 2019. "Autoencoder asset pricing models," *Yale ICF Working Paper No. 2019-04; Chicago Booth Research Paper No. 19-24*. <https://ssrn.com/abstract=3335536>.
- Jondeau, Eric, Poon, Ser-Huang, & Rockinger, Michael. 2007. *Financial modeling under non-gaussian distributions*. Springer Finance. Springer.
- Ke, Guolin, et al. 2017. "Lightgbm: A highly efficient gradient boosting decision tree," *Page 31493157 of: Proceedings of the 31st International Conference on Neural Information Processing Systems*. NIPS17. Red Hook, NY, USA: Curran Associates Inc.
- Kondratyev, Alexei. 2018. "Learning curve dynamics with artificial neural networks," *SSRN*. <https://ssrn.com/abstract=3041232>.
- Kritzman, Mark, Li, Yuanzhen, Page, Sébastien, & Rigobon, Roberto. 2011. "Principal components as a measure of systemic risk," *The Journal of Portfolio Management*, **37**(4), 112–126.
- Liu, Yan, & Zhang, Yi. 2018. "Low-dose CT restoration via stacked sparse denoising autoencoders," *Neurocomputing*, **284**, 80 – 89.
- Loretan, Mico, & English, William B. 2000. "Evaluating correlation breakdowns during periods of market volatility," *Board of Governors of the Federal Reserve System International Finance Working Paper*.
- Noureldin, Diaa, Shephard, Neil, & Sheppard, Kevin. 2012. "Multivariate high-frequency-based volatility (heavy) models," *Journal of Applied Econometrics*, **27**(6), 907–933.
- Packham, N., & Woebbecking, C.F. 2019. "A factor-model approach for correlation scenarios and correlation stress testing," *Journal of Banking and Finance*, **101**, 92 – 103.
- Pedregosa, F., et al. 2011. "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research*, **12**, 2825–2830.
- Preis, T, Kenett, DY, Stanley, HE, Helbing, D, & Ben-Jacob, E. 2012. "Quantifying the behavior of stock correlations under market stress," *Scientific Reports*, **752**(2).
- Rezek, I. A., & Roberts, S. J. 1998. "Stochastic complexity measures for physiological signal analysis," *IEEE Transactions on Biomedical Engineering*, **45**(9), 1186–1191.
- Shah, Nauman, & Roberts, Stephen. 2013. "Dynamically measuring statistical dependencies in multivariate financial time series using independent component analysis," *International Scholarly Research Notices*.
- Srivastava, Nitish, Hinton, Geoffrey, Krizhevsky, Alex, Sutskever, Ilya, & Salakhutdinov, Ruslan. 2014. "Dropout: A simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, **15**, 1929–1958.
- WRDS. 2019. *The Center for Research in Security Prices (CRSP)*

Index History - Intraday. <https://wrds-www.wharton.upenn.edu/>.
Zheng, Zeyu, Podobnik, Boris, Feng, Ling, & Li, Baowen. 2012.
“Changes in cross-correlations as an indicator for systemic risk.”,
Scientific Reports, **888**(2).

APPENDIX

A. Additional Training Details

Python Libraries: Deep sparse denoising autoencoders are defined and trained using the TensorFlow (Abadi *et al.*, 2015). For Gradient Boosted Decision Trees, we use the LightGBM library (Ke *et al.*, 2017) – using the standard LightGBMRegressor and LightGBMClassifier depending on the forecasting problem. The remainder of the models are implemented using standard scikit-learn classes (Pedregosa *et al.*, 2011) – with classes described in the hyperparameter optimisation section.

Hyperparameter Optimisation Details: Random search is conducted by sampling over a discrete set of values for each hyperparameter, which are listed below for each hyperparameter. For ease of reference, hyperparameters for all scikit-learn and LightGBM classes are referred to by the default argument names used in their respective libraries.

Deep Sparse Denoising Autoencoder

- Dropout Rate – [0.0, 0.2, 0.4, 0.6, 0.8]
- Regularisation Weight α – [0.0, 0.01, 0.1, 1.0, 10]
- Minibatch Size – [256, 512, 1024, 2048]
- Learning Rate – [10^{-5} , 10^{-4} , 10^{-3} , 10^{-2} , 10^{-1} , 1.0]
- Max. Gradient Norm – [10^{-4} , 10^{-3} , 10^{-2} , 10^{-1} , 1.0, 10.0]

Linear Regression

- Package Name – sklearn.linear_model
- Class Name – LogisticRegression
- 'alpha' – [10^{-5} , 10^{-4} , 10^{-3} , 10^{-2} , 10^{-1} , 1, 10, 10^2],
- 'fit_intercept' – [False, True],

Logistic Regression

- Package Name – sklearn.linear_model
- Class Name – LogisticRegression
- 'penalty' – ['l1']
- 'C' – [0.01, 0.1, 1.0, 10, 100]
- 'fit_intercept' – [False]

- 'solver' – ['liblinear']

Gradient Boosted Decision Tree

- Package Name – lightgbm
- Class Name – LGBMRegressor or LGBMClassifier
- 'learning_rate' – [10^{-4} , 10^{-3} , 10^{-2} , 10^{-1}]
- 'n_estimators' – [5, 10, 20, 40, 80, 160, 320]
- 'num_leaves' – [5, 10, 20, 40, 80]
- 'n_jobs' – [5]
- 'reg_alpha' – [0, 10^{-4} , 10^{-3} , 10^{-2} , 10^{-1}]
- 'reg_beta' – [0, 10^{-4} , 10^{-3} , 10^{-2} , 10^{-1}]
- 'boosting_type' – ['gbdt']

Multi-layer Perceptron

- Package Name – sklearn.neural_network
- Class Name – MLPRegressor or MLPClassifier
- 'hidden_layer_sizes' – [5, 10, 20, 40, 80, 160]
- 'activation' – ['relu']
- 'alpha' – [0, 10^{-4} , 10^{-3} , 10^{-2} , 10^{-1} , 1, 10, 10^2]
- 'learning_rate_init' – [10^{-4} , 10^{-3} , 10^{-2} , 10^{-1}]
- 'early_stopping' – [True],
- 'max_iter' – [500]