ARTICLE IN PRESS

Finance Research Letters xxx (xxxx) xxx

Contents lists available at ScienceDirect

Finance Research Letters

journal homepage: www.elsevier.com/locate/frl

Portfolio allocation under asymmetric dependence in asset returns using local Gaussian correlations

Anders D. Sleire [b], Bård Støve [b,*], Håkon Otneim [a], Geir Drage Berentsen [a], Dag Tjøstheim [b], Sverre Hauso Haugen [b]

[a] Department of Business and Management Science, Norwegian School of Economics, Norway
[b] Department of Mathematics, University of Bergen, Norway

ARTICLE INFO

Keywords:
Asymmetric dependence
Local Gaussian correlation
Mean–variance
Portfolio allocation

ABSTRACT

It is well known that there are asymmetric dependence structures between financial returns. This paper describes a portfolio selection method rooted in the classical mean–variance framework that incorporates such asymmetric dependency structures using a nonparametric measure of local dependence, the local Gaussian correlation (LGC). It is shown that the portfolio optimization process for financial returns with asymmetric dependence structures is straightforward using local covariance matrices. The new method is shown to outperform the equally weighted ("1/N") portfolio and the classical Markowitz portfolio when applied to historical data on six assets.

1. Introduction

Modern portfolio theory aims to allocate assets by maximizing the expected return while minimizing risk. Markowitz (1952) provides the foundation for the mean–variance (MV) approach under the crucial assumption that the asset returns follow a joint-Gaussian distribution. The idea is simple; highly correlated assets should be avoided to obtain a diverse portfolio. Several empirical studies, however, document asymmetries in the distribution of financial returns. In particular, one often observes stronger dependence between assets when the market is going down. This phenomenon is known as *asymmetric dependence structures*, see e.g. Silvapulle and Granger (2001), Campbell et al. (2002), Okimoto (2008), Ang and Chen (2002), Hong et al. (2007), Chollete et al. (2009), Aas et al. (2009), Garcia and Tsafack (2011).

Asymmetric dependence may lead to less effective diversification of mean–variance balanced portfolios. Several studies seek to overcome this shortcoming by modeling the dependence structure using copula theory and then applying this modeling into the portfolio allocation problem, see e.g. Patton (2004), Hatherley and Alcock (2007), Low et al. (2013), Kakouris and Rustem (2014), Bekiros et al. (2015) and Han et al. (2017). These procedures are quite complicated, and a non-technical asset manager might be overwhelmed by such choices. Moreover, there is no guarantee that portfolio allocations based on complex models will improve performance compared with simpler methods, see e.g., DeMiguel et al. (2009) and Low et al. (2016), who show that outperforming the naive $1/N$ portfolio remains an elusive task.

Without making assumptions about the nature of the underlying probability model, we present a simple adjustment to the MV approach by replacing the correlation matrix of the assets with a *local* correlation matrix. This approach is based on the local Gaussian correlation (LGC), see Tjøstheim and Hufthammer (2013), and has been applied successfully to analyze dependence structures between asset returns, see e.g., Støve and Tjøstheim (2014), Støve et al. (2014), Bampinas and Panagiotidis (2017) and Nguyen

* Corresponding author.
E-mail address: Bard.Stove@math.uib.no (B. Støve).

https://doi.org/10.1016/j.frl.2021.102475
Received 3 July 2021; Received in revised form 15 September 2021; Accepted 19 September 2021
Available online 25 September 2021
1544-6123/

Please cite this article as: Anders D. Sleire, *Finance Research Letters*, https://doi.org/10.1016/j.frl.2021.102475

et al. (2020), but has yet to be utilized in the portfolio allocation setting. The LGC provides a market-dependent adjustment to the correlation matrix that takes the current state of the market into account, and the main goal of this paper is to extend the classical MV framework by using the theory of the local Gaussian correlation, hence taking into account any asymmetric dependence structures between returns, and providing a simple alternative to the copula-based approaches. The organization of the paper is as follows. Section 2 presents the classical mean–variance portfolio approach and the extension using the local Gaussian correlation. In Section 3, we present a data set consisting of six asset returns, and in Section 4, we investigate the performance of portfolios constructed by our new approach and other methods. Finally, Section 5 offers some conclusions. In the supplementary material, we briefly review the local Gaussian correlation and provide some additional empirical study results.

2. Portfolio allocation using local Gaussian correlation

Denote the returns on N risky assets at time $t = 1, \dots, n$ by $\boldsymbol{R}_t \in \mathbb{R}^N$. Let $f_t(\boldsymbol{r}_t)$ denote the probability density function of $\boldsymbol{R}_t$, and let $\boldsymbol{\mu}_t \in \mathbb{R}^N$ and $\boldsymbol{\Sigma}_t \in \mathbb{R}^N \times \mathbb{R}^N$ denote its expectation vector and covariance matrix, respectively. Finally, let $\boldsymbol{w}_t = (w_{1,t}, \dots, w_{N,t})^T \in \mathbb{R}^N$ be the vector of portfolio weights at time t, to be determined by the portfolio selection strategy. We adopt the *full investment* constraint: $w_{1,t} + \cdots + w_{N,t} = 1$, for $t = 1, \dots, n$. Moreover, in the empirical example in Section 4, we investigate the performance of our proposed procedure with and without the *long-only* constraint ($0 \le w_{i,t} \le 1, i = 1, \dots, N, t = 1, \dots n$). We do not include a risk-free asset in our treatment of the portfolio allocation problem, but this does not impact our main findings.

The general portfolio optimization problem requires the investor to select weights $\boldsymbol{w}_t$ that maximizes an expected utility function at each time t. We consider the classical mean–variance as well as the minimum-variance utility functions, given respectively as

$$U_1 = \boldsymbol{w}_t^T \boldsymbol{\mu}_t - \frac{\gamma}{2} \boldsymbol{w}_t^T \boldsymbol{\Sigma}_t \boldsymbol{w}_t \qquad \text{and} \qquad U_2 = -\boldsymbol{w}_t^T \boldsymbol{\Sigma}_t \boldsymbol{w}_t, \tag{2.1}$$

where γ represents the investor's degree of risk aversion. Maximizing U_1 with respect to the portfolio weights provides a trade-off between expected volatility and expected returns for a given level of risk aversion (for simplicity fixed at $\gamma = 1$ throughout the paper). Maximizing U_2 results in the minimum variance portfolio.

We take as our point of departure the portfolio allocation approach as described by DeMiguel et al. (2009), Tu and Zhou (2011), and Low et al. (2016) when estimating the expected return vector $\boldsymbol{\mu}_t$ and covariance matrix $\boldsymbol{\Sigma}_t$ for monthly data. The approach is given as follows;

1. Select a sampling window of M trading months.
2. In each month $t > M$, estimate the expected return vector $\boldsymbol{\mu}_t$ and the covariance matrix $\boldsymbol{\Sigma}_t$ by their empirical counterparts, using the M preceding monthly returns.
3. Rebalance the portfolio on the first trading day of each month by solving the relevant optimization problem, i.e. optimizing one of the utility functions in (2.1).

The above algorithm implicitly assumes that the covariance matrix $\boldsymbol{\Sigma}_t$ completely describes the dependence structure among the assets under consideration. This property is not true in general unless the returns are jointly normally distributed, which is a strong assumption that is rarely satisfied in practice. Indeed, as mentioned in Section 1, there have been many attempts to replace the normality assumption in portfolio selection with more sophisticated distributions that better fit the return density f_t. However, this also results in a more complicated optimization routine than under the classical Markowitz framework indicated in the three steps listed above.

We propose to describe asymmetric dependence by making adjustments directly to the covariance matrix $\boldsymbol{\Sigma}_t$, which allows us to use the classical Markowitz formulation. To this end, we employ the *local Gaussian correlation*. The idea originated in Tjøstheim and Hufthammer (2013), who in turn based themselves on the local parameter concept of Hjort and Jones (1996). The latter authors approximate an unknown density function $f(\boldsymbol{r})$ by fitting a parametric family $f(\boldsymbol{r}, \theta)$ *locally* to $f(\boldsymbol{r})$, where $\theta \in \Theta$ is an unknown parameter in a parameter space Θ. This means that instead of constructing a single estimate $\hat{\theta}$ of θ, they rather estimate a parameter *function*, $\hat{\theta}(\boldsymbol{r})$, meaning that different members of the parametric family $\{f(\boldsymbol{r}, \theta), \theta \in \Theta\}$ approximate $f(\boldsymbol{r})$ in different parts of the domain of $f(\boldsymbol{r})$. Here, $\boldsymbol{r}$ represents a generic location in the domain of $\boldsymbol{R}_t$.

Hjort and Jones (1996) estimate $\theta(\boldsymbol{r})$ using a nonparametric local likelihood procedure, and Tjøstheim and Hufthammer (2013) consider the special case where $\{f(\boldsymbol{r}, \theta)\}$ is the multivariate Gaussian distribution, that is, $\theta = (\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Under this specification, it is natural to interpret the local covariance matrix $\boldsymbol{\Sigma}(\boldsymbol{r})$ as a measure of local dependence, which in particular gives a natural description of the asymmetric dependence relationships so often observed in financial returns.

Consider, for example, the observed returns on two of the assets in our data set displayed in the left panel of Fig. 1. The classical Gaussian assumption results in a single estimated covariance matrix $\hat{\boldsymbol{\Sigma}}$; while the local likelihood estimate $\hat{\boldsymbol{\Sigma}}(\boldsymbol{r})$ is a function of $\boldsymbol{r}$. In the right-hand panel of Fig. 1 we see the corresponding local Gaussian correlation, which clearly indicates that these returns are most strongly dependent in the lower left part of the distribution.

In order to incorporate the asymmetry observed in the right-hand panel of Fig. 1, we propose to replace step 2 in the above procedure with the following:

2.' In each month $t > M$, estimate the expected return vector $\boldsymbol{\mu}_t$ by its empirical counterpart, and the local Gaussian covariance matrix $\boldsymbol{\Sigma}_t(\boldsymbol{r})$, using the M preceding monthly returns.

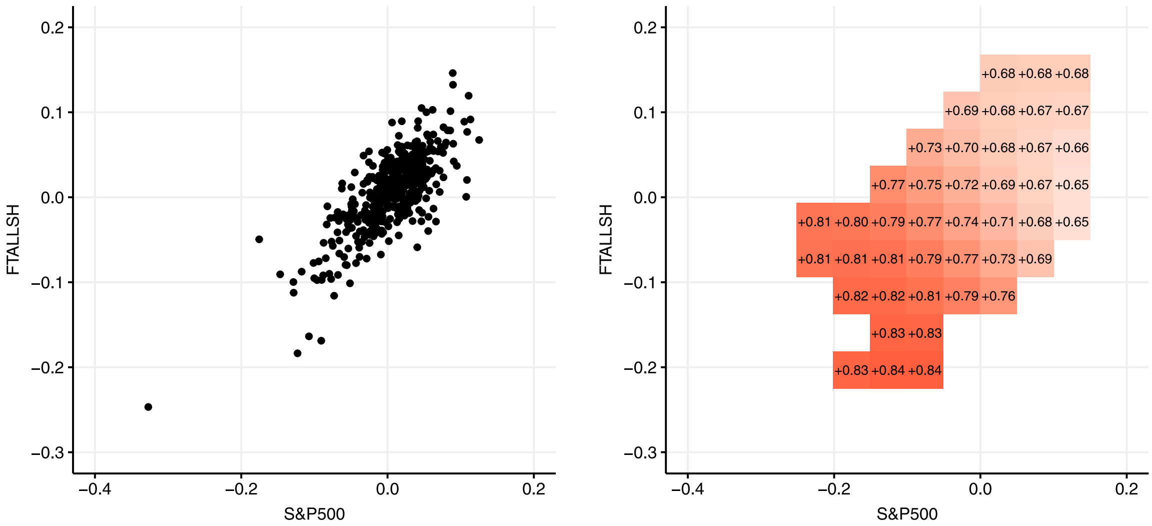

Fig. 1. Observations on two of the assets considered in the empirical analysis and the corresponding estimated local Gaussian correlation.

Table 1
Overview of the data series and abbreviations.

Name	Description
FTALLSH	FTSE Actuaries All Share Index
S&P500	Standard and Poor's 500 Index
BMUK10Y	UK Benchmark 10 Year DS government bond Index
BMUS10Y	US Benchmark 10 Year DS government bond Index
EWCI	Thomson Reuters Equal Weight Commodity Index
GSGCSPT	Standard and Poor's GSCI Gold Index

There are several approaches for choosing the evaluation point $\boldsymbol{r}$. A risk-averse investor can guard against large losses by selecting an evaluation point representing the asset returns during crisis periods. In this way, the corresponding estimated local covariance matrix reflects the (historical) dependence structure during crisis periods, provided that the window length M is chosen sufficiently large. However, the selection of the evaluation point can also be dynamic, i.e. $\boldsymbol{r} = \boldsymbol{r}_t$. For instance, the evaluation point may correspond to a subjective opinion of where the investor thinks the market is heading in the following trading month. The selection may also be based on more advanced statistical predictions.

In the empirical analysis in Section 4, we opt for a simple data-driven selection of evaluation points by computing the average of the last three months of observed returns. More specifically, the evaluation point at time t is defined for all pairs of assets i, j as

$$\boldsymbol{r}_t = \left(\frac{1}{3} \sum_{k=1}^{3} R^i_{t-k}, \frac{1}{3} \sum_{k=1}^{3} R^j_{t-k} \right). \tag{2.2}$$

This is a simple way of letting the covariance matrix dynamically adapt to the dependence structure of the market under the naïve assumption that the dependence structure between asset returns in month t is similar to the dependence structure of asset returns in the neighborhood of $\boldsymbol{r}_t$. As the empirical analysis in Section 4 will demonstrate, this simple selection of evaluation points performs well in practice.

3. Data

Our data set consists of monthly closing prices on six US dollar-denominated indices sourced from Refinitiv (past Thompson Reuters) Datastream. We calculate the returns as 100 times the difference in the log of the price indices. The sample period extends from February 1980 to August 2018, yielding 463 monthly return observations of the following assets: FTSE Actuaries All Share Index, (FTALLSH), Standard and Poor's 500 Index (S&P500), UK Benchmark 10 Year DS government bond Index (BMUK10Y), US Benchmark 10 Year DS government bond Index (BMUS10Y), Thomson Reuters Equal Weight Commodity Index (EWCI), and Standard and Poor's GSCI Gold Index (GSGCSPT) (see Table 1).

From the descriptive statistics in Table 2, we note that all of the returns are skewed and show relatively high kurtosis. Normality is rejected with the Jarque–Bera test on the 1% level for all series. This suggests that the multivariate normal distribution with a global covariance matrix does not provide a sufficient description of the dependence structure, particularly in the tails of the distribution.

Table 2
Correlations and descriptive statistics for the asset returns studied in the empirical analysis.

	FTALLSH	S&P500	BMUK10Y	BMUS10Y	EWCI	GSGCSPT
Global correlation matrix						
FTALLSH	1					
S&P500	0.760	1				
BMUK10Y	0.184	0.017	1			
BMUS10Y	−0.067	−0.029	0.489	1		
EWCI	0.246	0.288	−0.094	−0.185	1	
GSGCSPT	0.038	0.031	0.080	0.077	0.483	1
Local correlation matrix, bear market (lower 5% percentiles)						
FTALLSH	1					
S&P500	0.843	1				
BMUK10Y	0.174	−0.017	1			
BMUS10Y	0.020	0.034	0.635	1		
EWCI	0.161	0.185	−0.140	−0.224	1	
GSGCSPT	−0.135	−0.131	0.204	0.215	0.480	1
Descriptive statistics						
Observations	463	463	463	463	463	463
Mean	0.628	0.704	0.769	0.583	0.079	0.177
Std. Dev.	4.588	4.406	2.376	2.417	3.511	5.211
Variance	21.050	19.413	5.643	5.839	12.326	27.159
Skewness	−1.300	−0.968	−0.128	0.453	−0.592	0.026
Kurtosis	6.288	3.665	1.325	1.960	3.775	3.036
Jarque–Bera	903.903	335.969	36.135	91.622	306.377	180.971
Sharpe ratio	0.137	0.160	0.324	0.241	0.023	0.034
Max. drawdown	49.887	59.811	15.764	12.035	48.397	73.680
Min	−32.711	−24.677	−7.824	−7.600	−20.050	−21.887
1 Quartile	−1.474	−1.694	−0.585	−0.922	−1.794	−2.668
Median	1.176	1.242	0.843	0.497	0.151	−0.161
3 Quartile	3.559	3.265	2.151	1.853	1.998	2.899
Max	12.523	14.612	8.851	12.660	13.384	26.336

The two top panels in Table 2 show the global and local correlation matrices over the entire sampling period. The latter is constructed for a *bear market* scenario using the lower 5% percentiles for the evaluation point selection in the pairwise calculation approach described in the supplementary material. The strongest positive and negative correlation is observed between the stock indices FTALLSH and S&P500 ($\hat{\rho} = 0.76$), and between EWCI and BMUS10Y ($\hat{\rho} = -0.185$), respectively. The corresponding LGCs in the bear market scenario are $\hat{\rho} = 0.843$ and $\hat{\rho} = -0.224$, respectively, indicating the ability of the LGC to capture asymmetric dependence structures, see also Tjøstheim and Hufthammer (2013).

4. Empirical results

Our analysis[1] compares the portfolio selection strategies listed in Table 3 by evaluating their performance using both terminal wealth as well as a range of risk-adjusted performance measures. Following Low et al. (2016), we use the naïve $1/N$ weighted portfolio strategy as a benchmark model in the analysis. This strategy distributes weights equally across the portfolio at the start of the sampling period and is left unadjusted for the rest of the investment horizon. Tu and Zhou (2011) find that longer sampling windows result in improved portfolio strategy performance; hence we use both $M = 120$ and $M = 240$ month sampling windows. We report the $M = 240$ in the following sections. The corresponding results using $M = 120$ months are given in the supplementary material.

Inspired by Low et al. (2013), we proceed to evaluate the portfolio rebalancing, terminal wealth as well as the risk-adjusted performance for each of the strategies. A descriptive analysis of out-of-sample results is available in the supplementary material.

4.1. Portfolio rebalancing and terminal wealth

Table 4 provides a summary of the portfolio rebalancing analysis and the terminal wealth reached by each of the strategies. The average standard deviation within target portfolio weights across the entire out-of-sample time period is calculated as follows:

$$\bar{\sigma}_k = \frac{\sum_{t=1}^{n-M} \sigma_{t,k}}{n-M},$$

where

$$\sigma_{t,k} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(\hat{w}_{i,t,k} - \bar{w}_{\cdot,t,k})^2},$$

[1] Reproduce results or perform new studies with: https://gitlab.com/sleire/lgportf

Table 3
Portfolio strategies applied in the empirical analysis. All strategies allowing short sales have a lower limit on portfolio weights equal to −0.5.

Strategy	Description
Benchmark strategy	
EW	$1/N$ (equal weight portfolio) without rebalancing
Global approach	
MVS	Mean–variance portfolio with short sales
MVSC	Mean–variance portfolio with short sales constraint
MIN	Minimum variance portfolio
MINC	Minimum variance portfolio with short sales constraint
Local approach	
MVS-L	Mean–variance portfolio with short sales using local covariance matrices
MVSC-L	Mean–variance portfolio with short sales constraint using local covariance matrices
MIN-L	Minimum variance portfolio using local covariance matrices
MINC-L	Minimum variance portfolio with short sales constraint using local covariance matrices

Table 4
Portfolio rebalancing analysis and terminal wealth based on an initial investment of $1 for the different strategies (cf. Table 3) considered. Window size $M = 240$ months.

	$\bar{\sigma}_k$	Max. adj.	Min. adj.	Avg.turnover	Wealth	Wealth incl.tcost
Benchmark strategy						
EW	0	0	0	0	2.231	2.231
Global approach						
MVS	16.663	9.228	−6.741	6.325	2.605	2.551
MVSC	15.813	9.228	−7.426	5.622	2.571	2.524
MIN	16.249	5.036	−4.434	3.090	2.729	2.701
MINC	15.952	5.039	−4.428	2.793	2.742	2.717
Local approach						
MVS-L	15.660	73.184	−89.316	16.967	2.855	2.698
MVSC-L	14.253	18.096	−26.245	12.105	2.745	2.637
MIN-L	17.014	53.636	−79.931	21.617	2.953	2.749
MINC-L	15.243	19.646	−26.028	13.762	2.910	2.780

The Max.adj. and Min.adj. is the maximum values for positive and negative weight adjustments, respectively. Avg.turnover is defined in Eq. (4.1). Basis points of 15 per transaction are imposed as costs.

and where $\hat{w}_{i,t,k}$ is the portfolio weight for asset i at time t using portfolio strategy k, and $\bar{w}_{\cdot,t,k}$ is the average weight across the N assets in portfolio k at time t. The maximum values for positive and negative weight adjustments are the largest positive and negative weight changes on the asset level. Following DeMiguel et al. (2009), we also report the average turnover, which is calculated as

$$\text{Average turnover} = \frac{1}{n-M} \sum_{t=1}^{n-M} \sum_{i=1}^{N} (|\hat{w}_{i,t+1,k} - \hat{w}_{i,t,k}|), \tag{4.1}$$

We compute terminal wealth with and without a transaction cost of 15 basis points; such cost is comparable to prior studies, see e.g Low et al. (2016).

The variability of portfolio weights reported in Table 4 shows no systematic differences between the local and global approaches. Looking at the maximum and minimum adjustments of portfolio weights, however, there are clear differences. The local Gaussian strategies require adjustments of larger magnitude in both directions. This is particularly the case for the unconstrained models allowing short sales. Viewed across all strategies, we see larger and more frequent adjustments in the local portfolio strategies.

We see that increased trading volume translates into lower terminal wealth when transaction costs are included in the analysis. However, all local Gaussian strategies achieve higher terminal wealth than their traditional counterparts, also when transaction costs are included. The top-ranked strategy exclusive costs is MIN-L. When costs are included, the long-only portfolio MINC-L achieves the best result.

Fig. 2 shows wealth accumulation and drawdowns for the hypothetical investment of $1 in each of the nine strategies included in the analysis. As seen in the upper part of the figure, the local Gaussian MIN-L produces the largest final wealth when disregarding trade costs. It remains top-ranked during most months in the sample and suffers from smaller drawdowns in volatile periods such as the 2008 Financial Crisis. During this period, the EW strategy loses out substantially, which partially explains the overall poor performance of this strategy relative to the rest. The other local strategies also performs better than the corresponding global ones during this period. This is as expected, as we typically observe higher local dependence between asset returns during crisis periods (see e.g. Støve and Tjøstheim (2014)). When transaction costs are included, the strategy MIN-L still performs well but is surpassed by the constrained MINC-L, which has a lower turnover.

In Fig. 3, we see an illustration of the difference between the traditional Markowitz minimum variance, short sale constrained portfolio (MINC), and the proposed local counterpart (MINC-L). In the top panel, we see the estimated global and local variances of one of the assets in our example (EWCI) and the corresponding weight as a function of time. The local estimates are naturally

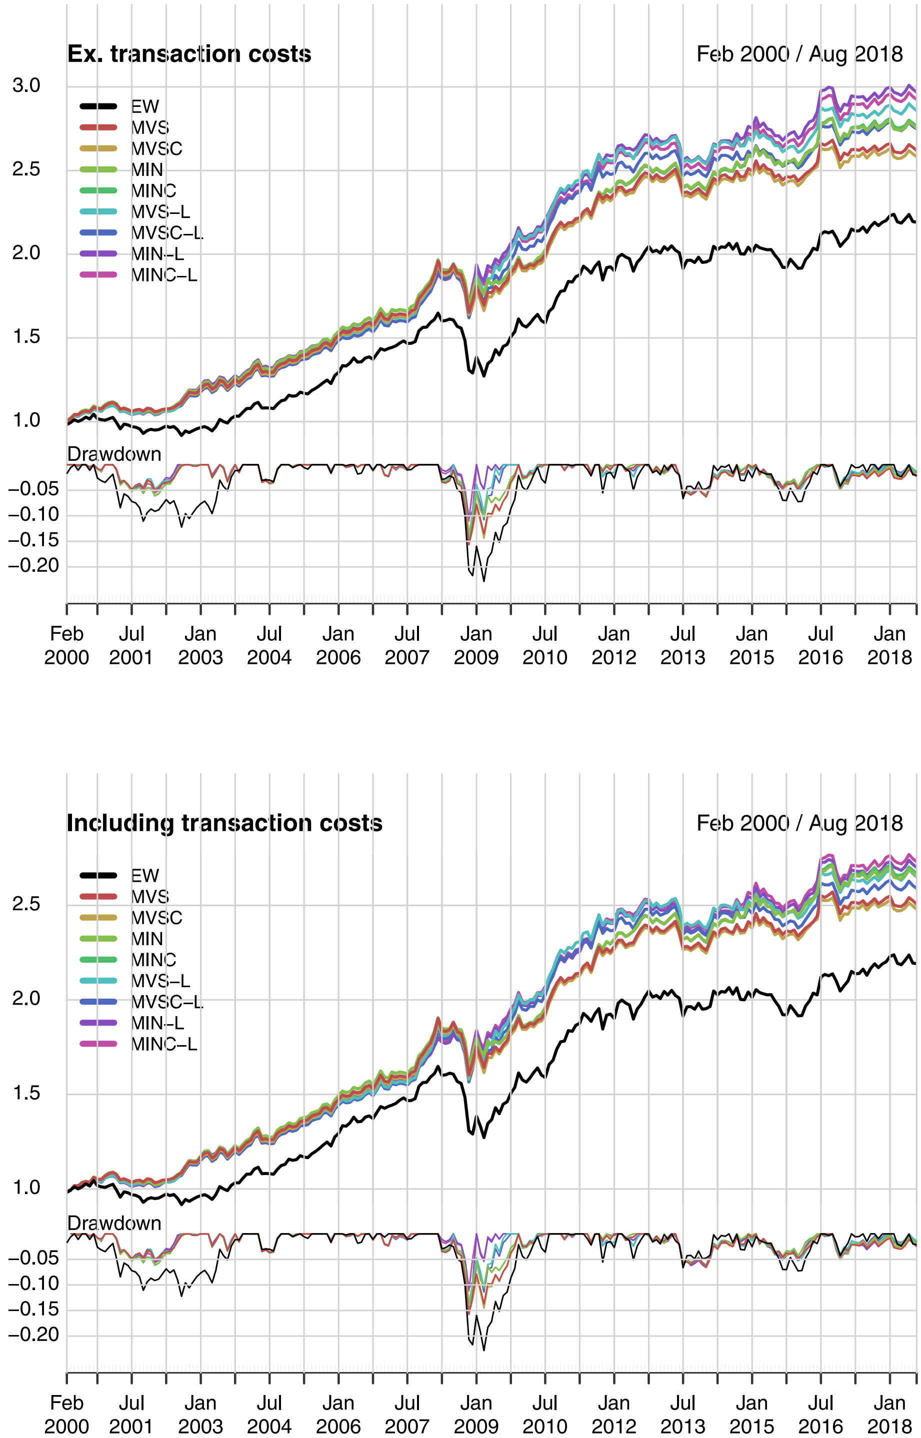

Fig. 2. Wealth accumulation of the different strategies based on an initial investment of $ 1, using a rolling window of size $M = 240$ monthly observations, top plot excluding transaction costs, bottom plot including transaction cost of 15 basis points per transaction.

more volatile due to their nonparametric estimation. However, they are also more sensitive to the state of the market, which is most easily visible during the financial crisis of 2008. The local variance of the asset increases sharply in this period, which is immediately reflected in the portfolio weight that quickly decreases to zero. We see similar effects for other instances of increased

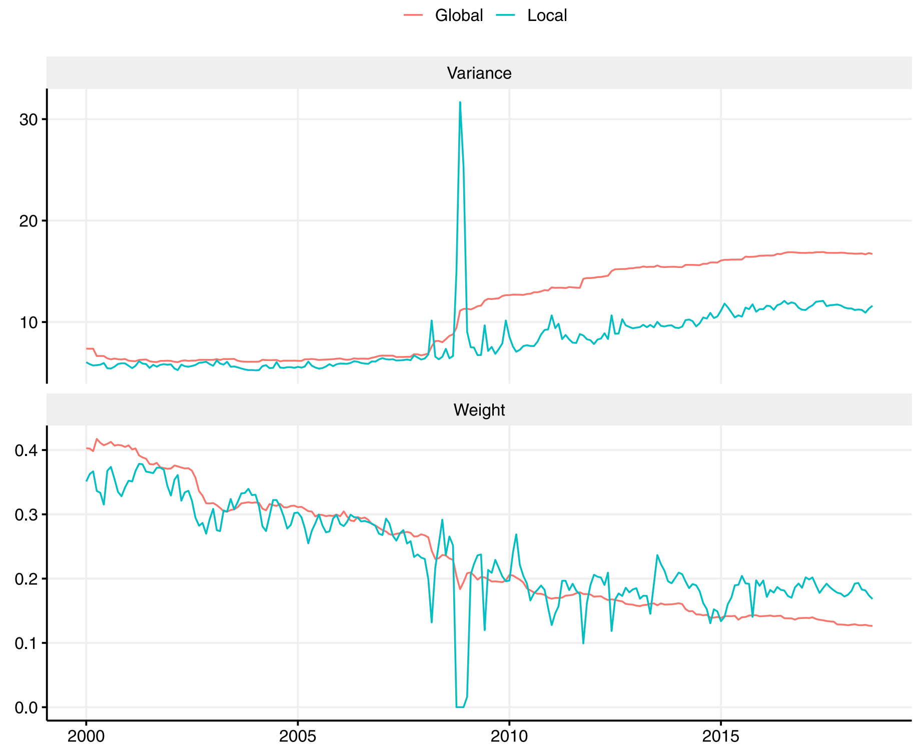

Fig. 3. Top panel: The estimated variance for one of the assets (EWCI) in our sample for the minimum variance and short sales constrained portfolio with traditional (global) Markowitz optimization (MINC), as well as the corresponding local version (MINC-L). In the bottom panel, we see the corresponding weight for this asset under the two strategies. See the online appendix for a corresponding plot covering all the assets in this example. The sample window is $M = 240$ months.

local variance. Note also that under the traditional (global) estimate, the estimated global variance reflects the financial crisis in the entire remaining sample period, leading to smaller investments in this asset than for the locally estimated portfolio.

4.2. Evaluation of risk-adjusted performance

Table 5 reports the out-of-sample performance of the different strategies considered using the following risk-adjusted metrics; The Sharpe ratio (Sharpe, 1966), the two modified Sharpe ratios VaR Sharpe, and ES Sharpe, where the Value at Risk and Expected Shortfall are used as risk measures Gregoriou and Gueyie (2003) and Favre and Galeano (2002). We also consider the Certainty Equivalent (CEQ), the Sortino ratio (Sortino and Price, 1994) and finally, the Omega ratio (Keating and Shadwick, 2002). All metrics produce high values for the best-performing strategies. Furthermore, we have performed the z-test of Ledoit and Wolf (2008), which is applied to the Sharpe Ratios to indicate the statistical differences for all MV optimizations against the $1/N$ benchmark.

Results excluding and including transaction costs are reported in Panel A and B, respectively. In both cases, the local portfolio strategies systematically outperform their traditional counterparts. Furthermore, we note that the local minimum variance portfolio has the highest performance across all metrics. The Sharpe and the annualized Sharp ratios prefer the long-only version (MINC-L), while the Var Sharpe, ES Sharpe, Sortino, and Omega ratios prefer the unconstrained version (MIN-L). The CEQ prefers the constrained version when costs are excluded and the unconstrained version when costs are included.

5. Concluding remarks

The results in this paper suggest that challenges related to return asymmetries may be handled in a familiar and well-established framework for portfolio management by replacing the global covariance matrix with a local version. Improved performance and simplicity are some of the appeals with the local Gaussian approach to portfolio management, even when considering higher transaction costs due to increased rebalancing requirements. There are, however, matters to keep in mind when implementing the approach. The selection of evaluation points for calculating the pairwise local correlations will affect the local Gaussian covariance matrix. We have evaluated alternative approaches to the moving evaluation point selection without observing substantial changes in results and conclusions. Nevertheless, there is a variety of options and possibilities for this choice. A more thorough analysis of these effects is left for future studies.

Table 5
Out-of-sample performance for the different portfolio strategies (cf. Table 3) considered. The sample window is $M = 240$ months.

	Sharpe	VaR Sharpe	ES Sharpe	Ann. Sharpe	CEQ	Sortino	Omega
Panel A: Ex. transaction costs							
Benchmark strategy							
EW	0.174	0.110	0.057	0.577	0.353	0.263	1.593
Global approach							
MVS	0.267	0.179	0.099	0.919	0.433	0.427	2.025
MVSC	0.264	0.175	0.097	0.905	0.426	0.417	2.009
MIN	0.289*	0.198	0.116	0.999	0.455	0.474	2.120
MINC	0.290*	0.199	0.117	1.000	0.457	0.474	2.120
Local approach							
MVS-L	0.276*	0.203	0.115	0.951	0.472	0.472	2.126
MVSC-L	0.270*	0.193	0.107	0.929	0.454	0.453	2.075
MIN-L	0.301*	**0.317**	**0.317**	1.041	**0.489**	**0.567**	**2.303**
MINC-L	**0.309***	0.240	0.157	**1.070**	0.482	0.554	2.268
Panel B: Incl. transaction costs							
Benchmark strategy							
EW	0.179	0.113	0.058	0.593	0.362	0.270	1.613
Global approach							
MVS	0.261	0.174	0.097	0.895	0.423	0.415	1.989
MVSC	0.258	0.171	0.095	0.885	0.418	0.407	1.978
MIN	0.285*	0.194	0.114	0.982	0.449	0.465	2.095
MINC	0.285*	0.195	0.115	0.985	0.451	0.466	2.097
Local approach							
MVS-L	0.262	0.188	0.106	0.900	0.448	0.440	2.041
MVSC-L	0.260	0.183	0.101	0.890	0.438	0.430	2.014
MIN-L	0.281	**0.279**	**0.279**	0.969	**0.456**	0.518	2.174
MINC-L	**0.295***	0.225	0.148	**1.019**	0.461	**0.522**	**2.182**

The maximum values for the risk-adjusted performance metrics are in **bold**. 15 basis points per transaction is imposed as costs in Panel B. '*' indicates that the Sharpe ratio is statistically different on the 5 percent level from the benchmark (EW) strategy using the z-test of Ledoit and Wolf (2008).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work has been partly supported by the Finance Market Fund (Norway). We thank two reviewers and the Editor for suggestions leading to an improved paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.frl.2021.102475.

References

Aas, K., Czado, C., Frigessi, A., Bakken, H., 2009. Pair-copula constructions of multiple dependence. Insurance Math. Econom. 44, 182–198.
Ang, A., Chen, J., 2002. Asymmetric correlations of equity portfolios. J. Financ. Econ. 63 (3), 443–494.
Bampinas, G., Panagiotidis, T., 2017. Oil and stock markets before and after financial crises: A local Gaussian correlation approach. J. Futures Mark. 37 (12), 1179–1204.
Bekiros, S., Hernandez, J.A., Hammoudeh, S., Nguyen, D.K., 2015. Multivariate dependence risk and portfolio optimization: An application to mining stock portfolios. Resour. Policy 46, 1–11.
Campbell, R., Koedijk, K., Kofman, P., 2002. Increased correlation in bear markets. Financ. Anal. J. 58 (1), 87–94.
Chollete, L., Heinen, A., Valdesogo, A., 2009. Modeling international financial returns with a multivariate regime switching copula. J. Financ. Econom. 7, 437–480.
DeMiguel, V., Garlappi, L., Uppal, R., 2009. Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? Rev. Financ. Stud. 22 (5), 1915–1953.
Favre, L., Galeano, J.-A., 2002. Mean-modified value-at-risk optimization with hedge funds. J. Altern. Invest. 5 (2), 21–25.
Garcia, R., Tsafack, G., 2011. Dependence structure and extreme comovements in international equity and bond markets. J. Bank. Financ. 35, 1954–1970.
Gregoriou, G., Gueyie, J.-P., 2003. Risk-adjusted performance of funds of hedge funds using a modified sharpe ratio. J. Wealth Manag. 6 (3), 77–83.
Han, Y., Li, P., Xia, Y., 2017. Dynamic robust portfolio selection with copulas. Finance Res. Lett. 21, 190–200.
Hatherley, A., Alcock, J., 2007. Portfolio construction incorporating asymmetric dependence structures: a user's guide. Account. Finance 47 (3), 447–472.
Hjort, N.L., Jones, M.C., 1996. Locally parametric nonparametric density estimation. Ann. Statist. (ISSN: 0090-5364) 24 (4), 1619–1647.
Hong, Y., Tu, J., Zhou, G., 2007. Asymmetries in stock returns: Statistical tests and economic evaluation. Rev. Financ. Stud. 20 (5), 1547–1581.

Kakouris, I., Rustem, B., 2014. Robust portfolio optimization with copulas. European J. Oper. Res. 235 (1), 28–37.
Keating, C., Shadwick, W.F., 2002. A universal performance measure. J. Perform. Meas. 6 (3), 59–84.
Ledoit, O., Wolf, M., 2008. Robust performance hypothesis testing with the sharpe ratio. J. Empir. Financ. 15 (5), 850–859.
Low, R.K.Y., Alcock, J., Faff, R., Brailsford, T., 2013. Canonical vine copulas in the context of modern portfolio management: Are they worth it? J. Bank. Financ. 37 (8), 3085–3099.
Low, R.K.Y., Faff, R., Aas, K., 2016. Enhancing mean–variance portfolio selection by modeling distributional asymmetries. J. Econ. Bus. 85, 49–72.
Markowitz, H., 1952. Portfolio selection. J. Finance 7 (1), 77–91.
Nguyen, Q.N., Aboura, S., Chevallier, J., Lyuyuan, Z., Zhu, B., 2020. Local Gaussian correlations in financial and commodity markets. European J. Oper. Res..
Okimoto, T., 2008. New evidence of asymmetric dependence structures in international equity markets. J. Financ. Quant. Anal. 43 (3), 787–816.
Patton, A.J., 2004. On the out-of-sample importance of skewness and asymmetric dependence for asset allocation. J. Financ. Econom. 2 (1), 130–168.
Sharpe, W.F., 1966. Mutual fund performance. J. Bus. 39 (1), 119–138.
Silvapulle, P., Granger, C.W.J., 2001. Large returns, conditional correlation and portfolio diversification: a value-at-risk approach. Quant. Finance 1 (5), 542–551.
Sortino, F.A., Price, L.N., 1994. Performance measurement in a downside risk framework. J. Invest. 3 (3), 59–64.
Støve, B., Tjøstheim, D., 2014. Measuring asymmetries in financial returns: An empirical investigation using local Gaussian correlation. In: N. Haldrup, M.M., Saikkonen, P. (Eds.), Essays in Nonlinear Time Series Econometrics. Oxford University Press, Oxford, pp. 307–329.
Støve, B., Tjøstheim, D., Hufthammer, K.O., 2014. Using local Gaussian correlation in a nonlinear re-examination of financial contagion. J. Empir. Financ. 25, 785–801.
Tjøstheim, D., Hufthammer, K.O., 2013. Local Gaussian correlation: A new measure of dependence. J. Econometrics 172, 33–48.
Tu, J., Zhou, G., 2011. Markowitz meets talmud: A combination of sophisticated and naive diversification strategies. J. Financ. Econ. 99 (1), 204–215.