Journal of Applied Finance & Banking, Vol. 14, No. 6, 2024, 69-94
ISSN: 1792-6580 (print version), 1792-6599(online)
https://doi.org/10.47260/jafb/1465
Scientific Press International Limited

Stocks, Gold and Crude Oil: How Valuable are Volatility and Correlation Timing?

Paolo Zagaglia[1]

Abstract

We study the sources of portfolio returns under parameter uncertainty arising from allocation strategies based on holdings of the NASDAQ index, WTI crude oil and gold. We investigate the contribution of volatility and correlation forecasts for dynamic portfolio allocations in the model of Brandt, Goyal, Clara, and Stroud (2005). We estimate alternative forecasting models using Bayesian methods, and evaluate investor's utility under the Bayesian predictive density. We then compare the performance of different portfolio strategies through both Sharpe ratios and utility-based metrics.
We show that dynamic strategies based on timing volatilities and correlations can add positively to the economic gains generated by non-diversified portfolios involving holdings of either crude oil or gold only. Hence, the economic benefits generated by holding crude oil and gold in asset allocations with stocks arise from the predictability of their volatilities and correlations.

JEL classification numbers: C22, C52, G11, G23.
Keywords: Bayesian models, dynamic asset allocation, multivariate correlation models, crude oil, gold.

---

[1] Department of Cultural Heritage, Università di Bologna, Ravenna, Italy; Rimini Centre for Economic Analysis.

Article Info: Received: October 25, 2024. Revised: November 16, 2024.
Published online: December 3, 2024.

1. Introduction

Crude oil and gold are typically assigned a peculiar role from both a macroeconomic and a financial point of view. A significant share of inflation fluctuations is attributed to changes in crude oil prices (e.g., see Ha, Kose, Ohnsorge, and Yilmazkuday, 2024). It has also been suggested that crude oil may be considered more akin to a purely financial asset rather than a commodity, whose underlying fundamentals consist in the imbalance between demand and supply (e.g., see Adams and Gluck, 2015; Kolodziej, Kaufmann, Kulatilaka, Bicchetti, and Maystre, 2014). On the other hand, gold is pointed at as a 'safe haven' asset or hedge for U.S. Dollar fluctuations (e.g., see Baur and Lucey, 2010; Reboredo, 2013).

The literature provides empirical evidence about a nonlinear mechanism for information transmission and price discovery between crude oil and precious metal markets (see Bildirici and Turkmen, 2015; Kumar, 2017).[2] This relation takes the form of volatility spillovers (see Hammoudeh and Yuan, 2008; Behmiri and Manera, 2015). A channel of dependence exists between crude oil and gold that runs also through exchange rate fluctuations. Evidence of strong dependence between the prices for crude oil and gold on one hand, and the U.S. Dollar on the other has been documented, especially over periods of financial turmoil (see Bedoui, Braeik, Goutte, and Guesmi, 2018).[3]

A number of contributions argue that crude oil and gold prices are related to stock prices. Christoffersen and Pan (2018) show that oil price volatility is a predictor for stock market volatility. At the same time, Ding, Granger, and Engle (1993) finds that extreme movements in the stock market carries explanatory power for changes in crude oil prices. Ferrer, Shahzad, Lopez, and Jareno (2018) suggest that the relation between financial market and crude oil prices has deepened since the financial crisis of 2007 in the U.S. The empirical results of Mensi, Ziadat, Rababa'a, Vo, and Kang (2024) identify asymmetric spillovers towards crude oil and gold that arise from changes in stock market prices: during periods of large market downturn or expansion, the effects of shocks from the stock market have a larger magnitude than in other circumstances.

In this paper, we study the scope for crude oil and gold in optimal portfolio allocations with stocks. We focus on the implications of volatility and correlation forecasting for portfolio strategies that address the effects of alternative sources for parameter uncertainty. Like Della Corte, Sarno, and Tsiakas (2013) and Ravazzolo and Lombardi (2016), we start from a dynamic asset allocation model that delivers utility-maximizing portfolio weights by taking as input the density forecasts for conditional volatilities and correlations. This allows us to investigate the contribution of different forecasting models to the performance of optimal asset

[2] The dependence between crude oil and precious metals has been documented for several countries, including China (see Zhang and Tu, 2016) and India (see Jain and Ghosh, 2013).

[3] Sari, Hammoudeh, and Soytas (2010) show that a weak long-term relation exists between the prices of precious metals, crude oil and the U.S. Dollar-Euro exchange rate, with a strong feedback taking place over the short run.

allocations. In order to estimate these models, we use Bayesian methods that deliver posterior densities for the parameters. In other words, we account explicitly for parameter uncertainty in our asset allocation model.

Following Barberis (2000) and Kandel and Stambaugh (1996) among others, we evaluate the investor's utility function under the predictive posterior distribution, and solve for utility-maximizing portfolio weights. The economic implications of parameter uncertainty can then be quantified by comparing measures of portfolio performance under two different types of scenarios: one where the parameters of the forecasting models take a set of values that are drawn from their posterior distributions - i.e., a Bayesian approach -, and another whereby the parameters are assigned only the posterior-mean estimates as their 'true' value. While the former accounts for estimation risk - i.e. uncertainty about the true values of estimated parameters -, the latter does not. At the same time, when we compare results across different forecasting models, we account for the effects of model uncertainty.[4]

In order to compute utility-maximizing portfolios with stocks, gold and crude oil, we use the framework for dynamic asset allocation proposed by Brandt, Goyal, Clara, and Stroud (2005). Our forecasting models include the constant conditional correlation model by Bollerslev (1990), the dynamic conditional correlation model by Engle and Sheppard (2001), and the 'BEKK' by Engle and Kroner (1995).[5] The dataset includes daily observations from January 2 1990 until September 3 2024.[6] We estimate the forecasting models over the first subsample until December 29 2023, and use the resulting parameter densities to initialize the out-of-sample forecast application carried out over the second - remaining - part of the sample. Based on data from January 2 2024 until September 3 2024, we re-estimate each forecasting model on recursive windows and compute one-step-ahead forecasts: the estimation window is expanded by one observation at the time and forecasts are computed until the end of the sample. We measure portfolio performance both through standard return-based indicators - i.e. the Sharpe ratio -, and through measures based on utility functions. Following Fleming, Kirby, and Ostdiek (2001), we quantify 'switching costs' and 'certainty-equivalent' returns. The former indicates the share of returns that an investor would be willing to pay in order to switch between two portfolio allocations generated by different strategies. The latter consists in the return that makes an investor indifferent between holding a portfolio with risky assets and holding a benchmark portfolio that delivers a riskless return.

We obtain two main results. Consistently with what has been shown for foreign exchange rates Della Corte, Sarno, and Tsiakas (2013), timing conditional volatilities and correlations generates positive gains in addition to strategies that

---

[4] Given the substantial computational burden that characterizes our optimal portfolio and estimation problems, we have chosen not to investigate the effects of model combination in this paper.

[5] Each model is estimated both on the Normal, and on the Student's *t* distribution.

[6] Our sample includes a number of episodes of boom-and-busts in financial markets. Episodes of instability can cause a breakdown of the correlation patterns that characterize non-turbulent periods (e.g., see Bernhart, Hocht, Neugebauer, Neumann, and Zagst, 2011). From a modelling perspective, market disruptions can exacerbate the problem of 'estimation risk'.

disregard parameter uncertainty, regardless of the sources of estimation risk. This means that there is predictability in the volatilities and correlations of assets whose prices have experienced low co-movements and periods with major spikes since 1990. However, not all the forecasting models are alike: in the case of our forecast-based rebalancing strategies with a stock market index, crude oil and gold, we document a clear ranking in the economic performance generated by the competing models. The second finding is that the portfolio holdings optimized for the three assets achieve higher economic gains with respect to those obtained from non-diversified portfolios allocated entirely to a single asset. Hence, the economic gains generated from holding crude oil and gold are based on the predictability of volatilities and correlations when parameter uncertainty is taken into consideration. In other words, the benefits from portfolio diversification concerns also the second moments of the distribution of asset returns.

The outline of the paper is as follows. Section 2 describes our framework for multivariate conditional modelling of returns. In section 3, we outline the dynamic setup for optimal asset allocation that we use in the empirical analysis. Section 4 discusses the dataset. Estimation results are presented in section 5. Section 6 focuses on performance of the optimized portfolios obtained from both in-sample and out-of-sample forecasting. Section 7 concludes the paper.

2. Multivariate correlation models

The starting point of our modelling approach to optimal portfolio allocation consists in the forecasts for returns, variances and correlations. We assume that the vector R_t of log-returns for stocks, bonds and crude oil at time t evolves according to

$$R_t=\Pi_t+\Omega_t^{1/2}\varepsilon_t \tag{1}$$

where Ω denotes the conditional covariance matrix, and ε_t is a vector of normally-distributed disturbances. In our empirical application, we make a simplifying assumption and impose a constant unconditional mean.

The first viable specification for the conditional covariance is one where $\Omega_t=\Omega$, which indicates the case of a multivariate linear regression model. A framework where the covariance changes over time with constant conditional correlations is the CCC model of Bollerslev (1990):

$$\Omega_t=\Sigma_t V \Sigma_t \tag{2}$$

$$\Sigma_t=diag(\sigma_{1,t}^{1/2}, \sigma_{2,t}^{1/2}, \sigma_{3,t}^{1/2}) \tag{3}$$

The diagonal elements of Σ_t are estimated from univariate GARCH models of order (1,1).

The dynamic conditional correlation model (DCC) of Engle (2002) extends the CCC in the following way:

$$\Omega_t = \Sigma_t V_t \Sigma_t \tag{4}$$

$$\Sigma_t = diag(\sigma_{1,t}^{1/2}, \sigma_{2,t}^{1/2}, \sigma_{3,t}^{1/2}) \tag{5}$$

$$V_t = (I \odot Q_t^{1/2}) Q_t (I \odot Q_t^{1/2}) \tag{6}$$

$$Q_t = (1-\alpha-\beta)\overline{Q} + \alpha\varepsilon_{t-1}\varepsilon_{t-1}' + \beta Q_{t-1} \tag{7}$$

where $\odot$ denotes the Hadamard product, I is an identity matrix, $\overline{Q}$ is the unconditional correlation matrix of ε_t, α and β are parameters. Finally, the BEKK model - of order (K=1, p=1, q=1) - introduced by Engle and Kroner (1995) provides for a rich parametrization of the conditional covariance:

$$\Omega_t = C + A'R_{t-j}R'_{t-j}A + B'\Omega_{t-j}B \tag{8}$$

In order to ensure both stationarity and positive definiteness of the covariance matrix, we follow a variant introduced by Rast, Martin, Liu, and Williams (2022) with all the diagonal elements $diag(A) \geq 0$ and $diag(B \geq 0)$. This yields the so-called positive-definite BEKK model. Finally, in order to account for fat-tailed returns, we consider two distributional assumptions for ε_t, namely a Normal and a Student's *t*.

2.1 Estimation, model comparison and forecasting

We estimate 8 models using Bayesian methods. The algorithms involved are standard and are discussed at length by Rast, Martin, Liu, and Williams (2022), amongst others. Let us denote by $\Theta := (\mathrm{u}, h_1, \ldots h_t)$ the vector of unobserved quantities, with a vector u of unknown parameters. Given a posterior density $p(\Theta|R_t)$, a prior distribution $p(\Theta)$ and a likelihood $p(R_t|\Theta)$, Bayes' theorem prescribes that

$$p(\Theta|R_t) \propto p(\Theta) p(R_y|\Theta) \tag{9}$$

Computing the marginal posterior distribution of parameters $p(u|R)$, a multidimensional integration is needed as

$$p(u|R_t) = \int_{h_1} \ldots \int_{h_T} p(u, \ldots h_T) dh_T \ldots dh_1 \tag{10}$$

Simulation-based integration techniques are then used to deal with this integration problem. In particular, a Markov Chain Monte Carlo method is employed to simulate from a Markov chain with stationary distribution equal to the density $p(\Theta|R_t)$.

Let us now turn to the choice of prior distributions. Since we estimate the models on log-returns of the data, we set the prior for Π to a Normal distribution with mean and variance obtained from a multivariate linear regression on a constant. The prior for the degrees of freedom v of the multivariate student's t distribution is truncated normal with $\mu=0$, $\sigma=30$ and lower truncation point at 2. The GARCH parameters underlying the DCC models are truncated in order to induce stationarity. Hence, they are drawn from uniform distributions $U(0, \alpha<1)$ and $U(0, (1-\beta)>\alpha)$. In the cases of both unconstrained and positive-diagonal BEKK, the priors are uniformly distributed. Their lower and upper bounds impose a stationarity condition, such that the eigenvalues of

$$A \otimes A + B \otimes B \tag{11}$$

in modulus are lower than one. For the positive-diagonal BEKK, we have prior distributions $A \sim U(a_1,a_2)$ and $B \sim U(b_1,b_2)$ with the additional restriction $(a_1, a_2, b_1, b_2)>0$. The prior for the constant correlation $\bar{Q}$ of the DCC model is distributed according to Lewandowski, Kurowicka, and Joe (2009) - denoted as LKJ - with scalar parameter m equal to 1. Following standard practice, we break down the constant covariance of all the BEKK models into $C=SPS$, with P as a correlation matrix and S is a diagonal matrix with standard deviations (see Alvarez, Niemi, and Simpson, 2014). We assume that $P \sim LKG(m=1)$ and that the standard deviations in S are log-normally distributed with mean 0.1 and standard deviation 0.7.

Since we focus on providing an economic evaluation of dynamic asset allocation strategies, we choose to refrain from carrying out a proper out-of-sample forecast evaluation of the - non-nested - conditional correlation models. However, in order to gain insights about their statistical properties, we compare their fitting abilities by using the deviance information criterion of Spiegelhalter, Best, Carlin, and van der Linde (2002):

$$DIC=2E[D(\Theta_M)]-D(E[\Theta_M]) \tag{12}$$

where Θ_M is the set of parameters in model M and $D(\cdot)$ is the deviance function, defined as minus twice the log-likelihood function. However, since the DIC is subject to Monte-Carlo sampling error, we follow Fioruci, Ehlers, and Filho (2014) and compute also the DIC weights, where the DIC score of every model is compared against the score of the best model:

$$w_M \propto exp(-(DIC_M - DIC_B)/2) \tag{13}$$

3. Dynamic asset allocation framework

In this paper we apply the framework for dynamic portfolio allocation proposed by Brandt, Goyal, Clara, and Stroud (2005). We consider the portfolio choice of an investor that maximizes end-of-period utility by allocating assets to a portfolio of N risky assets and a risk-free rate. The investor's optimization problem takes the form

$$V_t(W_t; Z_t) = \max_{\{w_s\}_{s=t}^{T-1}} E_T[U(W_T)|Z_T]; \; W_{s+1}=w_s R_{p;s+1}; \; \forall s \geq t \tag{14}$$

The term $R_{p,s+1}=w'_s(R_{s+1}-R^f)+R^f$ denotes the gross portfolio return from period s to s+1, w_s is a vector of portfolio weights for the risky assets, and R^f is the risk-free rate. Given the intertemporal dimension of the asset allocation problem, the investor takes into account the fact that portfolio holdings will change at each future point in time depending both on the available wealth, and on the information set. This dynamic effect is controlled for by the value function $V_t(W_t,Z_t)$.

We assume that instantaneous utility carries constant relative risk aversion:

$$U(W_T)=\frac{W_T^{1-\lambda}}{1-\lambda} \tag{15}$$

with λ as the coefficient of relative risk aversion. Brandt, Goyal, Clara, and Stroud (2005) follow a dynamic-programming approach, which delivers the first-order condition:

$$E_t\left[\partial U\left(w'_s(R_{s+1}-R^f)+R^f\right)\psi_{t+1}(Z_{t+1})\left(R_{t+1}-R^f\right)\right]=0 \tag{16}$$

The discrete-time problem is solved using a simulation method. Overall, the computation of the optimal asset allocation involves the following steps:

(A) we simulate paths of returns using one-step ahead forecasts for conditional means and correlations from estimated econometrics models. At each point in time t, we draw randomly the shock ε_t from the distribution $N(\mathbf{0},I)$ - with I as the identity matrix - to generate n=1,…10000 simulations of returns from

$$R_t=\Pi_t+\Omega_t^{1/2}\varepsilon_t \tag{17}$$

for given out-of-sample forecasts for Σ_{ε_t} and h_t. Obtaining analytical expressions for the forecasts from the different models is rather straightforward, except for the case of the DCC model. Following Engle and Sheppard (2001), we use the approximation

$$E_t(V_{t+n}) = \sum_{i=0}^{n-2}\left[(1-\alpha-\beta)\bar{V}(\alpha+\beta)^i + (\alpha+\beta)^{n-i}V_{n+1}\right] \tag{18}$$

The resulting information set of the investor consists of a stream of returns and state variables at time *t* for each path *s*. We can denote this set in a compact form as $I_t^s=\{R_t^s,Z_t^s\}$, where *s* denotes simulated values;

(B) we compute end-of-period time-T wealth for each *n*-path and solve the allocation problem backwards: at each *t* for every simulated path *n*, we use the recursion of the Bellman equation to solve for the optimal portfolio weights;

(C) we initialize the solution by computing the portfolio weights for the second-order Taylor approximation of the value function. This is then used as a starting guess for the solution of a fourth-order approximation of the value funtion;

(D) the process is iterated until the change in the optimal weights is numerically negligible.

Finally, in the evaluation of expected utility, we use the predictive distribution:

$$p\left(R_{t+1},\Pi_{t+1},h_{t+1},\Sigma_{t+1}|Z_t\right)=p(R_{t+1}|\Pi_{t+1},h_{t+1},\Sigma_{t+1},Z_t)p(\Pi_{t+1},h_{t+1},\Sigma_{t+1}|Z_t) \tag{19}$$

This means that the predictive distribution can change over time when the conditioning set Z_t incorporates additional data, which is factored into the posterior beliefs. As a result, the model accounts for Bayesian learning.[7]

[7] As pointed out by Della Corte, Sarno, and Tsiakas (2013), the Bayesian estimation framework provides the tools that are needed to consider two different dimensions of the dynamic asset allocation problem. The first one concerns the role of parameter uncertainty, while the second has to do with the predictability of means, variances and correlations. With both parameter uncertainty and predictability for time-varying volatilities and correlations, the optimal portfolio allocations account for both Bayesian learning and estimation risk. Under no predictability, there model contemplates only constant means, volatilities and correlations. With no parameter uncertainty, only the point estimates are factored into the solution of the dynamic planning problem.

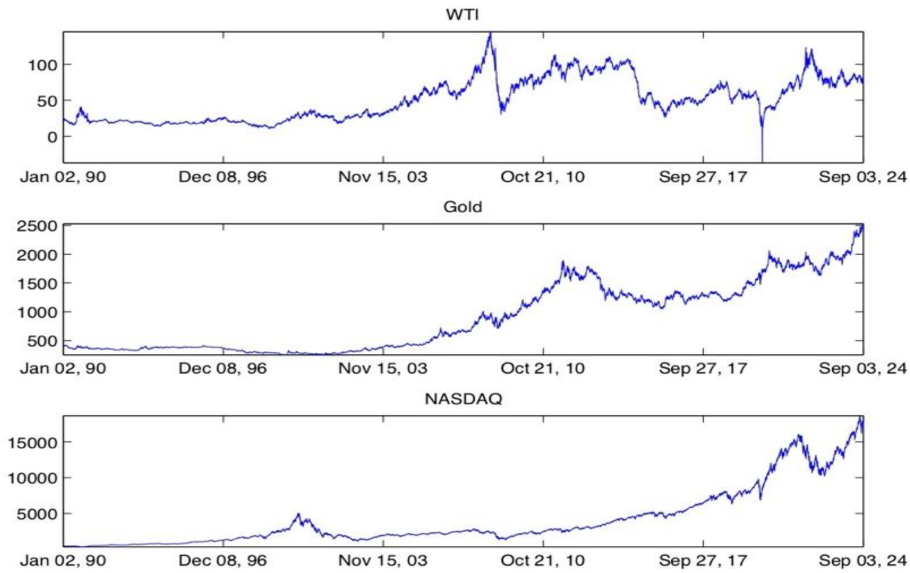

(a) Prices

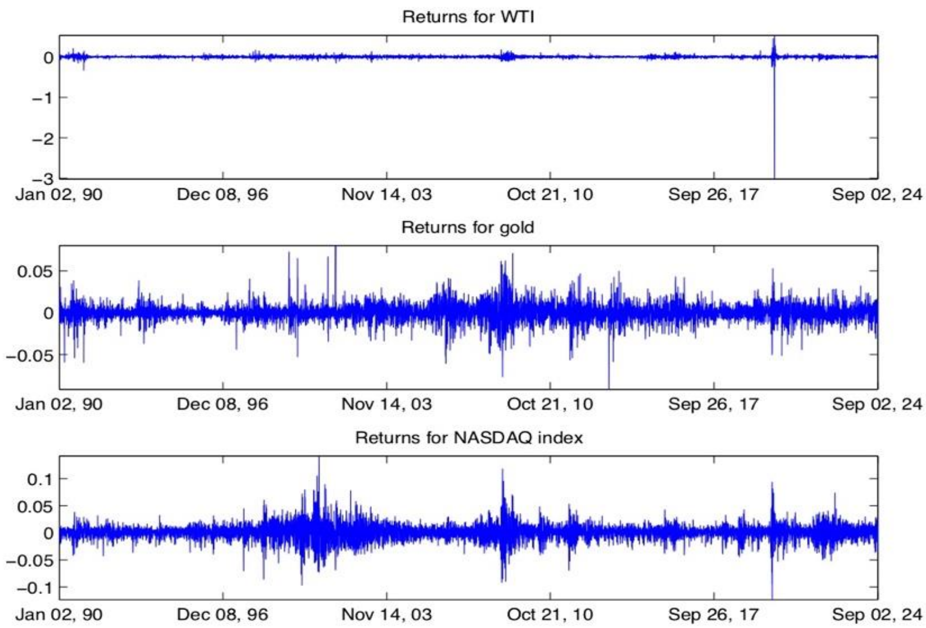

(b) Net returns

Figure 1: Plots of the daily dataset

4. Dataset and model calibration

We extract data from the FRED Database of the St Louis Federal Reserve Bank for West Texas Intermediate (Cushing, Oklahoma) spot prices and the NASDAQ composite index.[8] We also obtain old spot prices from the World Gold Council.[9] The risk-free rate is measured as the rate on a 3-month U.S. Treasury Bill, which is also retrieved from FRED. The full sample runs from January 2 1990 to September 3 2024. After removing missing values, the dataset includes 8702 observations, which are plotted in Figure 1.

Table 1 reports the descriptive statistics of net returns. The empirical distributions are left-skewed and peaked. The null of normality is rejected by the Jarque-Bera test statistics, especially for WTI returns. Table 1 includes also the statistics for the normality test of Anderson and Darling (1952), which confirms the rejection of the null of normality. The significance of the Ljung-Box statistics provides support for ARCH-type dynamics in the returns. These empirical facts are largely consistent with most studies on financial markets at the daily frequency. Interestingly, Spearman's rank correlation coefficients show that the returns for NASDAQ, spot gold and spot crude oil are characterized only by minor comovements. From an intuitive point of view, this should call into question the capability of extracting portfolio gains from forecasting the correlation between these assets.[10]

In order to compute optimized portfolios, we need to pin down two additional assumptions. We need to choose a value for the risk-aversion parameter λ of the utility function 15. Empirical studies based on microeconomic data tend to deliver estimates of λ in the range 1 to 4. Other studies such as Kandel and Stambaugh (1991) and Mehra and Prescott (1985) have proposed far higher value. We follow Barberis (2000) and calibrate λ to 5. The second assumption is related to the choice of the riskless asset, which is a 3-month U.S. Treasury Bill traded on the secondary market.

---

[8] The series can be downloaded, respectively, from https://fred.stlouisfed.org/series/DCOILWTICO and https://fred.stlouisfed.org/series/NASDAQCOM.

[9] Available for download from https://www.gold.org/goldhub/data/gold-prices.

[10] It should be mentioned that linear correlations provide a similar picture.

Table 1: Descriptive statistics for net returns

Panel (a): Summary statistics			
	WTI	Treasury bond	Nasdaq
Mean	0.00008	0.0001	0.0005
Std dev	0.044	0.021	0.015
Skewness	-38.767	-0.078	-0.024
Kurtosis	2571.107	9.692	9.756
Panel (b): Normality and ARCH tests			
Anderson-Darling	17.455 [0.08]	11.310 [0.4]	9.601 [0.03]
JB	1.7e+3*	2.1e+3*	1.9e+2*
LJB(12)	16.13*	21.04*	13.01*
Panel (c): Spearman's rank correlations			
	Brent crude	Treasury bond	Nasdaq
Brent crude	1		
Treasury bond	0.095	1	
Nasdaq	0.091	-0.020	1

Legend: Brackets report the marginal probability. The LJB(12) is the Ljung-Box test statistics on the squared residuals from the regression of the conditional mean. Under the null of no serial correlation, it is distributed as a $\chi^2(q)$ distribution with q lags and a criticical value equal to 21.03. JB is the Jarque-Bera test of normality. It has a χ^2 distribution with 2 degrees of freedom. The critical value at the 5% level is 5.99.

5. Estimation results

We start by estimating the models on a subsample with 8534 observations until December 29 2023. The posterior-mean estimates provide the starting values for a recursive estimation and out-of-sample forecast application over the second part of the sample. Running from January 2 2024 until September 2 2024, the estimation window is expanded by one observation at the time, the models are re-estimated, and one-step-ahead forecasts are computed.

The potential numerical issues that may arise in the - automatic - recursive estimation imposes a burden on the Markov Chain Monte Carlo algorithm to both run smoothly, and deliver well-behaved chains. However, this problem has already been tackled in a number of previous contributions. We follow the standard practice of relying on short pilot runs to find an initial guess for the parameter estimates, which are then used as a starting point for the full algorithm. The length of the pilot run is set to 10000 iterations for each model. The main algorithm is then run for 60,000 iterations, from which we disregard the first 10000 as burn-in. The simulated Markov chains are automatically assessed for convergence through the test diagnostic of Brooks and Gelman (1998). In addition, in the estimation of the models based on the first part of the sample, we also check the plots of the running means of the marginal posteriors as suggested by Bauwens, Lubrano, and Richard (1999), and apply the separated partial means tests of Geweke (2005, p. 149).

In order to provide some understanding on the results of the estimation from the first part of the sample, we can consider the in-sample statistical fit of the models. Table 2 reports the values for the DIC statistics obtained from estimation on data until December 29 2023. Based on the deviance information criterion, a clear ranking emerges among the models. The use of a Student's *t* distribution generates superior models with respect to a Normal distribution. In addition, regardless of any distributional assumption, BEKK models deliver improvements in fit with respect to both CCC and DCC models. Hence, allowing for a more flexible structure and a richer parametrization in the model for conditional correlations allows to capture features of the data that are otherwise harder to match under tighter modelling assumptions. On the other hand, imposing a positivity constraint on the diagonal ARCH and GARCH parameters in a BEKK model delivers a better fit. Despite these considerations, we should notice that the DIC statistics are somewhat similar across the model space: the pd-BEKK with *t* distribution has an improvement in DIC by 10% with respect to the CCC with a Normal distribution. This suggests that there can be a statistical case for ranking models through DIC weights, thus providing a clear-cut picture.

Table 2: DIC model comparison

Model	***DIC***	w_M
CCC - Normal	1851.20	0.0000
DCC - Normal	1820.06	0.0000
BEKK - Normal	1809.51	0.0000
pd-BEKK - Normal	1801.99	0.0000
CCC - *t*	1803.38	0.0000
DCC - *t*	1794.06	0.0000
BEKK - *t*	1755.82	0.0911
pd-BEKK - *t*	1751.22	0.9089

Since the *t*-distributed positive-diagonal BEKK model appears to deliver the best fit, and in the interest of parsimony, we discuss the key estimation results only for this model. Table 3 includes the point estimates, along with numerical standard deviations and 95% credible intervals estimated as 95% highest density intervals. All the numerical standard deviations are of a small magnitude, indicating that the parameter estimates have converged towards the posterior means. There is no difference in the precision with which the conditional mean parameters are estimated with respect to the parameters driving the conditional volatilities and correlations. The last row of Tables 3 indicates high values for the degrees of freedom of the *t* distribution. This indicates that a distribution with fat tails provides a reasonable assumption. Finally, there is evidence of persistence in both conditional volatilities and correlations.

Table 3: Point estimates and 95% credible intervals for the positive-diagonal BEKK model with *t* distribution

Parameter	Estimate	Standard deviation	Credible intervals
μ_1	0.00120	0.00044	(-0.0190, 0.0173)
μ_2	0.00137	0.00058	(-0.0164, 0.0180)
μ_3	0.00218	0.00071	(-0.0183, 0.0166)
A [1, 1]	0.2300	0.0710	(0.1791, 0.2150)
A [2, 2]	0.2390	0.0500	(0.1740, 0.2756)
A [3, 3]	0.1977	0.0118	(0.1522, 0.2300)
C [1, 1]	0.2509	0.0185	(0.1408, 0.3509)
C [2, 1]	0.3110	0.0170	(0.1701, 0.4308)
C [2, 2]	0.4071	0.0315	(0.2900, 0.6911)
C [2, 3]	0.2760	0.0509	(0.1371, 0.3119)
C [3, 1]	0.2049	0.0840	(0.0920, 0.3050)
C [3, 3]	0.4460	0.0790	(0.2700, 0.6202)
B [1, 1]	0.7103	0.0910	(0.7460, 0.8687)
B [2, 2]	0.7514	0.0580	(0.8248, 0.8671)
B [3, 3]	0.7855	0.0811	(0.8640, 0.8850)
v	6.1900	0.407	(1.0370, 2.9600)

Legend: the table reports credible intervals from the 95% highest posterior density region for each parameter estimate, which is the shortest interval containing 95% of the posterior distribution. Numerical standard errors are also included.

Panels (a) and (b) of Figure 2 plot the in-sample forecasts for conditional variances and correlations from the *t*-distributed positive-diagonal BEKK. The model captures well-known spikes in market turbulence, such as those that took place both in 2009 and in 2020 as a consequence of the Covid outbreak. Estimated positive correlations appear to have become stronger over time after 2010. The pd-BEKK with *t* distribution captures no visible break after the Lehman bankruptcy of 2008, which is often associated with the beginning of a period of positive correlations with increasing size. Out-of-sample forecasts can be found in panels (c) and (d). The *t*-distributed positive-diagonal BEKK picks up the increase in volatility that has characterized both oil and gold prices since June 2024. The resulting spike in stock market volatility goes hand-in-hand with an increase in correlation between gold and NASDAQ index prices on one hand, and a drop in the correlation between oil and NASDAQ.

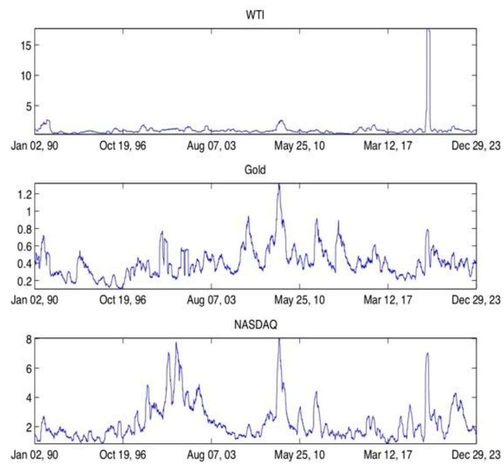

(a) In-sample variance forecasts

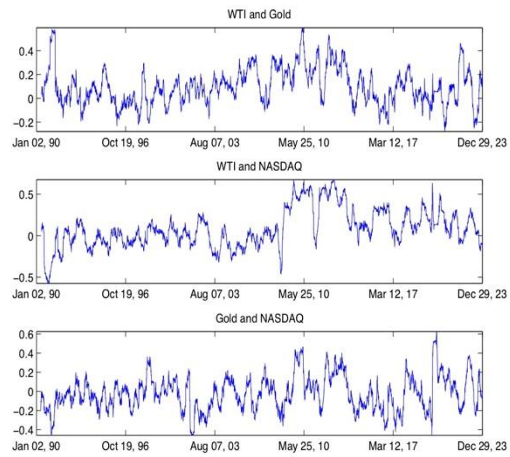

(b) In-sample correlation forecasts

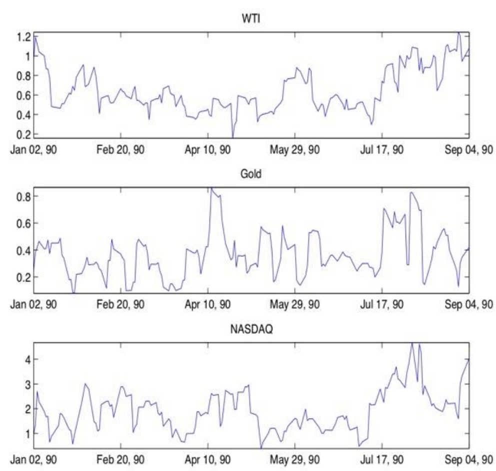

(c) Out-of-sample variance forecasts

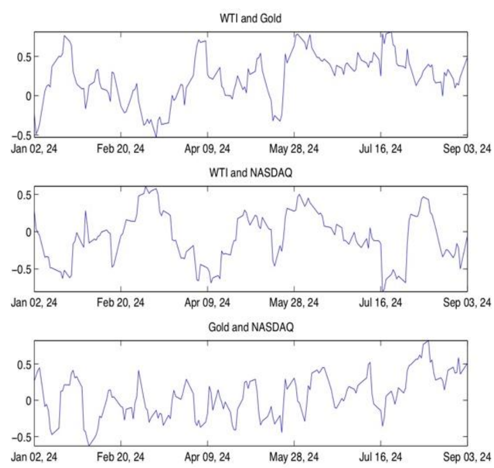

(d) Out-of-sample correlation forecasts

Figure 2: Forecasts of conditional variances and correlations for the t-distributed positive-diagonal BEKK

6. The economic value of optimal dynamic portfolios

The solution to the optimal portfolio allocation problem of the representative investor involves two stages. Based on the posterior-mean estimates obtained from data until December 29 2023, the investor starts by computing one day-ahead forecasts for returns, conditional volatilities and correlations. Then, the utility-maximization problem is solved by choosing optimal portfolio weights, thereby generating portfolio rebalancing strategies on a daily frequency. And a full cycle including recursive estimation, forecast computation, and utility maximization through portfolio rebalancing is repeated until September 2 2024.[11] Given the substantial computational costs involved in the recursive estimation and in the solution of the dynamic programming problem, we only consider daily rebalancing strategies that bear a transaction cost c equal to 5 basis points.

In order to provide evidence on the economic implications of volatility and correlation forecasting, we compute optimal 'static' portfolios based on constant variances and covariances. These are obtained from multivariate linear models of excess returns estimated with Bayesian methods. From an intuitive point of view, the static allocation can be interpreted as a baseline portfolio composition without rebalancing based on volatility and correlation forecasting.

Since the Bayesian estimation framework allows us to study the role of estimation risk, we compare the economic outcomes from portfolio allocations based on different degrees of parameter uncertainty. For this purpose, we consider a benchmark where any source of uncertainty is switched off. In a 'plug-in method' for assigning parameters, we fix parameter values to their estimates (see Della Corte, Sarno, and Tsiakas, 2013). Then, we move on to optimal portfolios in models with uncertain volatilities and correlations, where the unconditional mean is fixed to the estimates from a multivariate linear regression. Finally, we compare the results obtained earlier with the implications from parameter uncertainty for means, volatilities and correlations.

We are left with the problem of measuring the economic implications of our dynamic portfolio strategies. For the sake of comparison with the available literature, we compute Sharpe ratios based on ex-post annualized mean portfolio returns and annualized standard deviation. However, it should be kept in mind that Sharpe ratios are affected by a number of issues. As suggested by Han (2006), they disregard the role of higher moments, and are subject to various forms of distortions arising from dynamic allocation strategies. Hence, it becomes appealing to rely on metrics that have economic underpinnings. Following Fleming, Kirby, and Ostdiek (2001) and West, Edison, and Cho (1993) amongst others, we compute the fee Φ that an investor would be willing to pay in order to switch from portfolio a to portfolio b in order to achieve a similar average utility level. This can also be interpreted as a breakeven cost, above which switching across strategies is not desirable. Given

[11] As a result, also the estimates for the conditional-mean parameters change between periods when we run the full round of estimation and simulation to solve for the optimizing portfolio weights.

$R^{b}_{p,t+1}$ and $R^{a}_{p,t+1}$ as the returns from benchmark and alternative allocations, we compute Φ such that:

$$\sum_{t=0}^{T-1} E_t\left[U\left(R^{a}_{p,t+1}/exp(\Phi)\right)\right] = \sum_{t=0}^{T-1} E_t\left[U\left(R^{b}_{p,t+1}\right)\right] \tag{20}$$

We consider four static benchmark portfolios:

- an allocation 'R' based on the constant conditional covariance estimated from a multivariate regression model of excess returns;
- an allocation *S* with only stocks;
- an allocation *W* with only WTI crude oil;
- an allocation *G* with only gold.

6.1 Optimal in-sample portfolios

Table 4 reports the performance of optimal allocations based on in-sample forecasts. Panels (a) summarizes the key portfolio statistics where there is no estimation risk. Panel (b) shows the results with parameter uncertainty in volatilities and correlations, while panel (c) summarizes the portfolio statistics in models with full Bayesian learning. A number of interesting results arise, some of which are consistent with the findings on other assets obtained by Della Corte, Sarno, and Tsiakas (2013) and Ravazzolo and Lombardi (2016).

First of all, higher Sharpe ratios are generated as the sources of parameter uncertainty are factored into optimal portfolios allocations. Average return variance increases less than average excess returns when estimation risk becomes widespread. Disregarding estimation risk generates portfolios characterized by the lowest average returns and volatilities. On the other hand, addressing the effects of full uncertainty penalizes portfolio performance with respect to the case of uncertainty in volatilities and correlations. To put it differently, even in a portfolio of assets carrying low unconditional correlation, there are returns that can be generated by forecasting conditional volatilities and correlations, in addition to the returns derived from optimal allocations that disregard parameter uncertainty.

In order to interpret these findings, we should stress that there is a trade-off between the depth of estimation risk arising from model size, and the increase in predictive ability generate by larger models under parameter uncertainty. In our modelling framework, the latter overcomes the former. However, improvements in forecast performance depend on the source of parameter uncertainty. Our random walk model for returns is clearly unable to deliver a satisfactory forecast performance. And this can explain the worsening in portfolio statistics with full uncertainty with respect to uncertainty in volatilities and correlations.

Table 4: Performance of optimal in-sample portfolios

Model	μ_p	σ_p	SR	$\Phi(R)$	$\Phi(S)$	$\Phi(W)$	$\Phi(G)$
(a) Plug-in method							
MLR	18.700	12.109	1.016	-	-	-	-
CCC - Normal	22.800	12.614	1.300	490.0	511.0	501.0	301.4
DCC - Normal	26.140	13.720	1.438	539.0	546.7	535.7	337.9
BEKK - Normal	26.305	13.910	1.431	812.0	610.8	577.0	401.8
pd-BEKK - Normal	26.955	14.200	1.447	833.0	612.0	574.8	449.0
CCC - *t*	27.961	14.790	1.457	931.0	686.0	673.4	350.9
DCC - *t*	28.505	14.850	1.488	934.8	738.0	708.0	361.2
BEKK - *t*	28.977	15.290	1.476	1029.0	906.5	775.5	709.8
pd-BEKK - *t*	29.79 0	15.520	1.507	1127.0	896.7	802.6	720.3
(b) Uncertainty in volatilities and correlations							
MLR	18.499	11.903	1.008	-	-	-	-
CCC - Normal	25.553	13.395	1.429	502.7	526.3	536.7	387.3
DCC - Normal	25.957	13.662	1.431	553.0	563.1	542.6	417.7
BEKK - Normal	26.765	13.959	1.458	836.9	629.1	583.4	420.0
pd-BEKK - Normal	27.068	14.157	1.459	854.6	630.3	582.2	439.3
CCC - *t*	28.28	14.949	1.463	955.2	706.5	682.1	458.9
DCC - *t*	28.785	14.989	1.493	959.1	760.1	717.2	469.5
BEKK - *t*	29.088	14.751	1.538	1195.8	933.7	785.5	726.2
pd-BEKK - *t*	29.290	14.157	1.616	1256.3	923.6	863.6	806.5
(c) Uncertainty in means, volatilities and correlations							
MLR	18.008	11.582	1.002	-	-	-	-
CCC - Normal	25.809	13.596	1.427	500.1	520.0	524.4	372.5
DCC - Normal	26.217	13.867	1.429	544.1	556.9	540.0	409.4
BEKK - Normal	27.033	14.168	1.456	852.0	617.2	581.0	412.0
pd-BEKK - Normal	27.339	14.369	1.457	860.2	620.7	676.0	420.5
CCC - *t*	28.563	15.173	1.460	940.0	699.7	680.5	442.0
DCC - *t*	29.073	15.213	1.490	944.3	750.9	714.7	460.5
BEKK - *t*	29.379	14.972	1.534	1152.0	920.9	780.1	713.3
pd-BEKK - *t*	29.583	14.369	1.613	1226.0	916.0	849.1	790.4

Legend: μ_p and σ_p denote annualized average portfolio returns and volatilities. *SR* is the Sharpe ratio. $\Phi(\cdot)$ is the cost that an investor is willing to pay in order to switch from a benchmark to an alternative portfolio. The alternative portfolio is based on in-sample forecasts produced by each model reported in the model. Benchmark portfolios include only stocks (*S*), holding only WTI crude oil (*W*), only gold (*G*), and a portfolio of assets generated from a multivariate regression model of excess returns.

The only exception to this finding concerns the dynamic allocations under constant volatilities and correlations estimated from linear regressions, which suffer from a drop of Sharpe ratios as additional sources of parameter uncertainty are introduced. Following the reasoning discussed earlier, using a constant covariance matrix implies removing the sources for possible gains from predictability of volatilities and correlations.

Our results show that distributional assumptions for the forecasting models matter. Models with fat tails are associated with larger Sharpe ratios than models based on the Normal distribution. At the same time, imposing constant volatilities and correlation yields the most ‘cautious’ portfolios across different degrees of parameter uncertainty.

Forecasting conditional volatilities and correlations through multivariate models with a richer structure of interactions contribute to deliver higher Sharpe ratios. The family of BEKK models outperforms the in-sample capabilities of alternative models. On the other hand, we also find that imposing parametric restrictions that have economic underpinnings does deliver gains in portfolio performance: the positive-diagonal BEKK achieves consistently higher returns and lower volatility especially when parameter uncertainty is considered. This indicates that there can be benefits from responding to estimation risk when a model with a sufficiently flexible structure is adopted.

Let us now focus on the costs of switching across from static portfolios to dynamic allocations. We consider four benchmark - static - portfolios based on non-diversified holdings with the NASDAQ index, WTI crude oil, gold, as well as a constant-weight allocation obtained from a linear regression model. The first noticeable result is that the fees for switching towards forecast-based strategies have a positive sign for all the correlation models, regardless of the benchmark portfolio. Consistently with the results for the Sharpe ratios, the switching costs increase as parameter uncertainty is accounted for by optimal portfolio allocations: the improvements in average returns lead to higher utility levels that are discounted by the investor.

Portfolio holdings with only the NASDAQ index leads to the highest losses with respect to dynamic allocations. All in all, large opportunity costs can emerge both from disregarding uncertainty in volatilities and correlations, and from choosing portfolios with fixed asset allocations.

A benchmark allocation with only gold holdings generates the lowest breakeven costs against the alternative portfolios, while the benchmark linear-regression allocation produces consistently higher fees. This means that a static portfolio entirely allocated to gold generates economic gains that are closer to those obtained from the alternative portfolio strategies, with respect to the performance of other benchmarks. All these results suggest that the source of economic gains from holding gold or crude oil lies in the contributions that they deliver for portfolio diversification.

Table 5: Performance of optimal out-of-sample portfolios

Model	μ_p	σ_p	SR	$\Phi(R)$	$\Phi(S)$	$\Phi(W)$	$\Phi(G)$
(a) Plug-in method							
MLR	11.151	10.816	0.439	-	-	-	-
CCC - Normal	13.035	10.584	0.626	337.8	255.5	336.2	246.7
DCC - Normal	13.893	10.656	0.703	467.9	426.0	332.3	260.5
BEKK - Normal	15.354	11.448	0.782	495.2	330.8	340.5	258.1
pd-BEKK - Normal	15.134	10.872	0.803	511.6	473.5	394.4	346.1
CCC - *t*	15.961	11.736	0.814	540.1	304.8	440.8	270.5
DCC - *t*	16.448	12.168	0.825	552.8	305.7	523.1	278.5
BEKK - *t*	16.504	12.168	0.830	620.1	598.5	550.7	547.2
pd-BEKK - *t*	17.105	12.456	0.859	670.5	610.9	602.0	555.3
(b) Uncertainty in volatilities and correlations							
MLR	10.602	10.332	0.406	-	-	-	-
CCC - Normal	13.188	10.074	0.673	360.1	270.9	347.8	251.7
DCC - Normal	13.776	10.143	0.720	475.9	405.1	343.8	265.8
BEKK - Normal	14.897	10.688	0.710	510.1	345.7	352.2	263.4
pd-BEKK - Normal	15.408	10.971	0.821	524.0	494.8	407.9	353.2
CCC - *t*	16.276	11.316	0.872	527.5	318.5	456.0	275.9
DCC - *t*	16.877	11.592	0.903	530.8	359.4	541.2	284.1
BEKK - *t*	16.850	11.661	0.893	586.2	625.3	569.7	558.3
pd-BEKK - *t*	17.492	12.075	0.918	699.2	638.3	622.8	566.5
(c) Uncertainty in means, volatilities and correlations							
MLR	10.065	9.576	0.382	-	-	-	-
CCC - Normal	12.219	12.054	0.482	280.3	233.6	210.7	206.9
DCC - Normal	13.251	12.136	0.564	391.0	372.7	311.5	248.0
BEKK - Normal	14.706	12.382	0.670	477.6	429.5	380.2	316.5
pd-BEKK - Normal	14.967	13.038	0.657	455.5	414.3	369.7	290.3
CCC - *t*	15.481	13.366	0.679	490.3	436.7	413.2	226.9
DCC - *t*	16.035	13.858	0.695	504.9	330.5	490.4	233.6
BEKK - *t*	16.438	14.186	0.707	511.4	523.7	516.3	459.1
pd-BEKK - *t*	16.087	13.858	0.699	501.1	534.6	564.4	465.8

Legend: μ_p and σ_p denote annualized average portfolio returns and volatilities. *SR* is the Sharpe ratio. $\Phi(\cdot)$ is the cost that an investor is willing to pay in order to switch from a benchmark to an alternative portfolio. The alternative portfolio is based on in-sample forecasts produced by each model reported in the model. Benchmark portfolios include only stocks (*S*), holding only WTI crude oil (*W*), only gold (*G*), and a portfolio of assets generated from a multivariate regression model of excess returns.

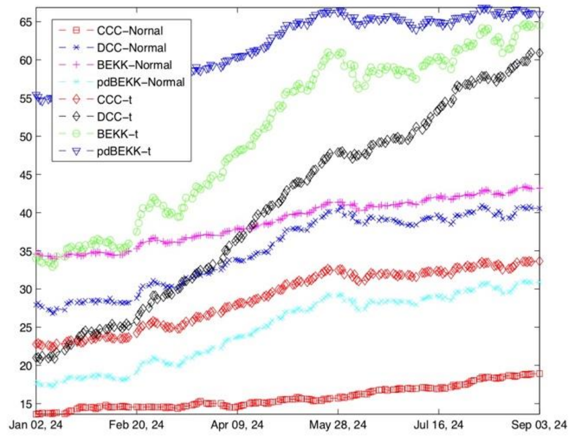

(a) Differential with respect to linear-regression allocation

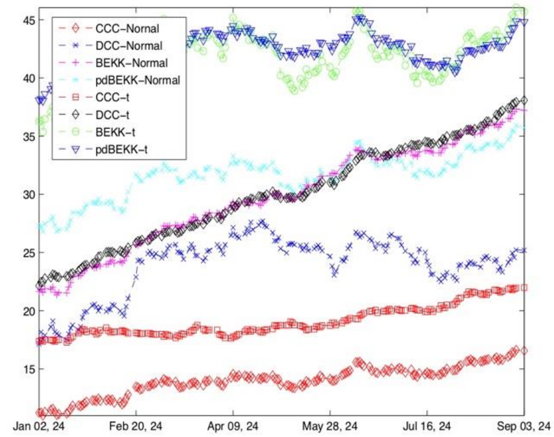

(b) Differential with respect to NASDAQ-only allocation

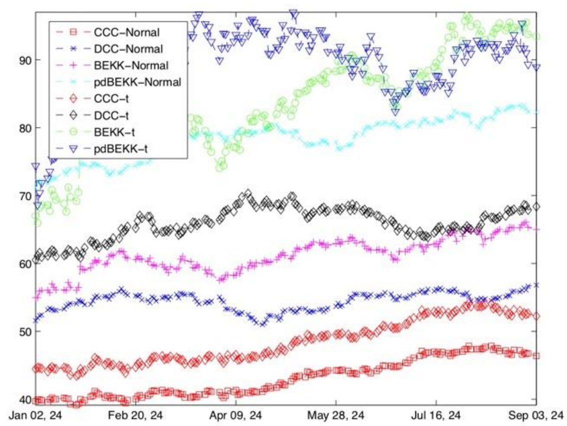

(c) Differential with respect to WTI-only allocation

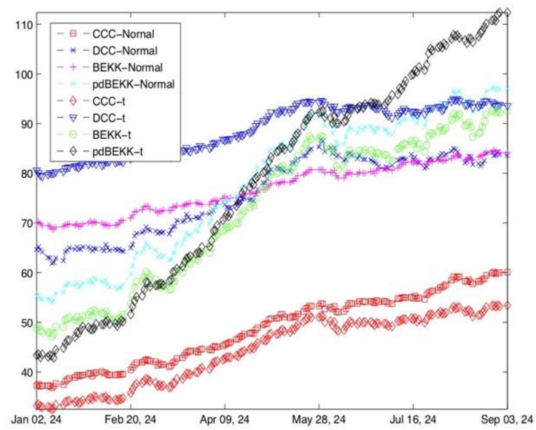

(d) Differential with respect to gold-only allocation

Figure 3: Certainty-equivalent return differentials

Legend: the figures report certainty equivalent return differentials over a one-day forecast horizon with respect to the certianty equivalent returns of alternative benchmark strategies.

6.2 Optimal out-of-sample portfolios

Let us now turn to optimal portfolio allocations evaluated out of sample. Table 5 reports the summary statistics for different degrees of parameter uncertainty. The drop in average portfolio returns is larger than the reduction of average volatility generated by all the models: this is the cost arising from the struggle of our predictive models to produce reliable forecasts when the information on which they are based does not cover the entire dataset. The rest of the patterns that characterize optimal in-sample portfolios emerge also in the case of out-of-sample dynamic allocations. Hence, volatility and correlation forecasts contribute positively both to return-based and to utility-based metrics of performance. However, the challenge to forecast the means of returns prevents full Bayesian learning from delivering additional benefits in favour of the portfolio allocations. While distributional assumptions do make a different, so does the choice of the dynamic correlation model.

In order to gain additional insights on the sources of economic benefits generated by the forecasting models, we evaluate the certainty-equivalent return associated to each portfolio. This is the riskless return that delivers the same utility level arising from an alternative portfolio of risky assets. In other words, that would make an investor indifferent between choosing an allocation of riskless assets and an allocation of risky assets. Given the utility function 15, we use the results obtained under uncertainty in volatilities and correlations to compute:

$$CER_{t+1}=U^{-1}\left(E_t\left(U(W_{t+1})\right)\right) \tag{21}$$

Rather than reporting end-of-period values, we study the evolution of certainty-equivalent returns in comparison to each benchmark static allocation introduced earlier for measuring the switching costs across strategies. The cumulative certainty-equivalent return over the out-of-sample forecast period can be written as:

$$CCER^{a}_{t+1}=\sum_{m=0}^{T-1}\left(CER^{a}_{m,t+1}-CER^{b}_{m,t+1}\right) \tag{22}$$

A cumulative certainty-equivalent return that increases over time indicates that the alternative model records a stronger return-generation performance than that of the benchmark allocation. Consistently with the results of Table 5, Figure 3 shows that all the forecast-based portfolios are capable of generating net economic gains over the forecast period. Even the constant-correlation models are capable of a positive performance with respect to each benchmark allocation. However, the growth rate of the cumulative certainty-equivalent returns differs across models. For instance, when compared against the static allocation with only the NASDAQ index, the positive-diagonal BEKK model with *t*-distributed disturbances typically starts from the highest values at the beginning of the forecast evaluation period, and then suffers from a number of periods with negative returns. Non-diversified portfolios with full allocations to either gold or crude oil perform very poorly with respect to dynamic utility-maximizing portfolios. This supports the finding discussed previously for which there are benefits from portfolio diversification also in out-of-sample allocation decisions.

7. Conclusion

The goal of this paper is to study the sources of economic value that can be extracted from forecast-based strategies for the construction of optimal portfolios with holdings of the NASDAQ index, WTI crude oil and gold. We use Bayesian methods to estimate alternative models for conditional volatilities and correlations. We then compute out-of-sample forecasts for conditional means, volatilities and correlations of the returns. The forecasts from each competing model are used to solve for optimal dynamic portfolios in the model of Brandt, Goyal, Clara, and Stroud (2005) where investor's utility is evaluated at the Bayesian predictive density. This way, we can factor different degrees of parameter uncertainty into the construction of portfolios. Optimal asset allocation are evaluated using standard financial metrics - such as the Sharpe ratio -, as well as criteria based on economic principles. We obtain two main results. First of all, timing conditional volatilities and correlations generates positive gains in addition to those delivered by strategies involving holdings of one asset only. Second, the sources of the economic benefits generated by either WTI crude oil or gold rely on their contributions to portfolio diversification.

Our analysis can be extended along several directions. Since a daily rebalancing frequency can have an effect on the ability to time volatilities and correlations, we should explicitly model the fact that the costs of changing asset weights affect the transition to a new portfolio. For instance, we could modify our asset allocation problem by introducing the idea of transaction cost aversion proposed by Collin-Dufresne, Daniel, Moallemi, and Saglam (2012). We should also examine the role of our assumption about the utility function of investors. First of all, the effects of changes of the coefficient of relative risk aversion should be assessed properly. We should also compare the performance of optimal portfolios under different functional forms of utility, like the logarithmic utility, or the function due to Epstein and Zin (1989). As stressed by Buraschi, Porchia, and Trojani (2010), forecasting correlations in the context of optimal portfolio allocation raises the issue of correlation risk. While our discussion provides insights on the implications of this source of risk for portfolio composition, we do not attempt to quantify directly the resulting demand for hedging.

From the point of view of statistical modelling, we should study the effects of model uncertainty on our optimal asset allocations. This would amount to extending our work by comparing the performance of alternative schemes for generating forecasts. Avramov (2002) suggests that model uncertainty can have portfolio implications more compelling than those of estimation risk. Disregarding model uncertainty may lead investors to favor portfolio strategies with poor performance. In our case, it would be interesting to apply the forecast-combination framework of Bates and Granger (1969). To be consistent with our estimation methods, the Bayesian Model Averaging approach presents appealing features. Finally, since our sample includes a period of substantial market instability, we should model explicitly the dynamics of higher moments. This would allow us to extend our results to the case of distributional timing discussed by Jondeau and Rockinger (2006) among others.

References

[1] Adams, Z. and Gluck, T. (2014). Financialization in commodity markets: A passing trend or the new normal? Journal of Banking and Finance, 60:93–111, 2015.

[2] Alvarez, I., Niemi, J. and Simpson, M. (2014). Bayesian inference for a covariance matrix. unpublished manuscript available at https://arxiv.org/pdf/1408.4050.

[3] Anderson, T. W., and Darling, D. A. (1952). Asymptotic theory of certain goodness-of-fit criteria based on stochastic processes. Annals of Mathematical Statistics, 23.

[4] Avramov, D. (2002). Stock return predictability and model uncertainty. Journal of Financial Economics, 64(3), pp. 423–458.

[5] Barberis, N. (2000). Investing for the long run when returns are predictable. Journal of Finance, 55(1), pp. 225–264.

[6] Bates, J. M. and Granger, C. W. J. (1969). The combination of forecasts. Operations Research Quarterly, 20(4), pp. 451–468.

[7] Baur, D. G. and Lucey, B.M. (2010). Is gold a hedge or a safe haven? an analysis of stocks, bonds and gold. Financial Review, 45(2), pp. 217–229.

[8] Bauwens, L., Lubrano, M. and Richard, J. F. (1999). Bayesian Inference in Dynamic Econometric Models. Oxford University Press.

[9] Bedoui, R., Braeik, S., Goutte, S. and Guesmi, K. (2018). On the study of conditional dependence structure between oil, gold and USD exchange rates. International Review of Financial Analysis, 59:134–146.

[10] Behmiri N. B., and Manera, M. (2015). The role of outliers and oil price shocks on volatility of metal prices. Resources Policy, 46(2), pp. 139–150.

[11] Bernhart, G., Hocht, S., Neugebauer, M., Neumann, M. and Zagst, R. (2011). Asset correlations in turbulent markets and the impact of different regimes on asset management. Asia-Pacific Journal of Operational Research, 28(1), pp. 1–23.

[12] Bildirici, M. E. and Turkmen, C. (2015). Nonlinear causality between oil and precious metals. Resources Policy, 46(2), pp. 202–211.

[13] Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: A multivariate generalized arch model. The Review of Economics and Statistics, 72(3), pp. 498-506.

[14] Brandt, M. W., Goyal, A., Santa Clara, P. and Stroud, J. R. (2005). A simulation approach to dynamic portfolio choice with an application to learning about return predictability. Review of Financial Studies, 18(2), pp. 831–873.

[15] Brooks, S. P. and Gelman, A. (1998). General methods for monitoring convergence of iterative simulations. Journal of Computational and Graphical Statistics, 7(4), pp. 434–455.

[16] Buraschi, A., Porchia, P. and Trojani, F. (2010). Correlation risk and optimal portfolio choice. Journal of Finance, 65(1), pp. 393–420.

[17] Christoffersen, P. and Pan, X. (2018). Oil volatility risk and expected stock returns. Journal of Banking and Finance, 95, pp. 5–26.

[18] Collin-Dufresne, P., Daniel, K., Moallemi, C. C. and Saglam, M. (2012). Strategic asset allocation with predictable returns and transaction costs. unpublished manuscript available at http://www.princeton.edu/~msaglam/dpc.pdf.

[19] Della Corte, P., Sarno, L. and Tsiakas, I (2013). Volatility and Correlation Timing in Active Currency Management, in Handbook of Exchange Rates, Chapter 15, pp. 221–263. Wiley, London.

[20] Ding, Z., Granger, C. W. J. and Engle, R. R. (1993). A long memory property of stock market returns and a new model. Journal of Empirical Finance, 1, pp. 83–106.

[21] Engle, R. F. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. Journal of Business and Economic Statistics, 20(3), pp. 339–350.

[22] Engle, R. F. and Kroner, K. (1995). Multivariate Simultaneous Generalized ARCH. Econometric Theory, 11, pp. 122–150.

[23] Engle, R. F. and Sheppard, K. (2001). Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH. NBER Working Paper 8554.

[24] Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. Econometrica, 57(4), pp. 937–969.

[25] Ferrer, R., Shahzad, S. J. H., Lopez, R. and Jareno, F. (2018). Time and frequency dynamics of connectedness between renewable energy stocks and crude oil prices. Energy Economics, 76, pp. 1–20.

[26] Fioruci, J. A., Ehlers, R. S. and Filho, M. G. A. (2014). Bayesian multivariate garch models with dynamic correlations and asymmetric error distributions. Journal of Applied Statistics, 41(2), pp. 320–331.

[27] Fleming, J., Kirby, C. and Ostdiek, B. (2001). The economic value of volatility timing. Journal of Finance, 56:329–352, 2001.

[28] Geweke, J. F. (2005). Contemporary Bayesian Econometrics and Statistics. John Wiley and Sons.

[29] Ha, J., Kose, M. A., Ohnsorge, F. and Yilmazkuday, H. (2024). What explains global inflation. IMF Economic Review.

[30] Hammoudeh, S. and Yuan, Y. (2008). Metal volatility in presence of oil and interest rate shocks. Energy Economics, 30(2), pp. 606–620.
[31] Han, Y. (2006). Asset allocation with a high-dimensional latent factor stochastic volatility model. Review of Financial Studies, 19(1), pp. 237–271.
[32] Jain, A. and Ghosh, S. (2013). Dynamics of global oil prices, exchange rate and precious metal prices in India. Resources Policy, 38(1), pp. 88–93.
[33] Jondeau, E. and Rockinger, M. (2006). The economic value of distributional timing. unpublished manuscript available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=963288.
[34] Kandel, S. and Stambaugh, R. F. (1991). Asset returns and intertemporal preference. Journal of Monetary Economics, 27(1), pp. 39–71.
[35] Kandel, S. and Stambaugh, R. F. (1996). On the predictability of stock returns: An asset allocation perspective. Journal of Finance, 51(2), pp. 385-424.
[36] Kolodziej, M., Kaufmann, R. K., Kulatilaka, N., Bicchetti, D. and Maystre, N. (2014). Crude oil: Commodity or financial asset? Energy Economics, 46(1), pp. 216–223.
[37] Kumar, S. (2017). On the nonlinear relation between crude oil and gold. Resources Policy, 51, pp. 219–224.
[38] Lewandowski, D., Kurowicka, D. and Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. Journal of Multivariate Analysis, 100(9), pp. 1989–2001.
[39] Mehra, R. and Prescott, E. C. (1985). The equity premium: A puzzle. Journal of Monetary Economics, 15(2), pp. 145–161.
[40] Mensi, W., Ziadat, S. A., Al Rababa'a, A. R., Vo, X. V. and Kang, S. H. (2024). Oil, gold and international stock markets: Extreme spillovers, connectedness and its determinants. The Quarterly Review of Economics and Finance, 95(C), pp. 1–17.
[41] Rast, P., Martin, S., Liu, S. and Williams, D. (2022). A new frontier for studying within-person variability: Bayesian multivariate generalized autoregressive conditional heteroskedasticity models. Psychol Methods, 27(5), pp. 856-873.
[42] Ravazzolo, F. and Lombardi, M. (2016). On the correlation between commodity and equity returns: Implications for portfolio allocation. Journal of Commodity Markets, 2(1), pp. 45–57.
[43] Reboredo, J. C. (2013). Is gold a safe haven or a hedge for the us dollar? implications for risk management. Journal of Banking and Finance, 37(3), pp. 2665–2676.
[44] Sari, R., Hammoudeh, S. and Soytas, U. (2010). Dynamics of oil price, precious metal prices, and exchange rate. Energy Economics, 32(2), pp. 351–362.
[45] Spiegelhalter, D. J., Best, N. G., Carlin, B. P. and van der Linde, A. (2002). Bayesian measures of model complexity and fit. Journal of Royal Statistical Society, Series B, 64(4), pp. 583–639.

[46] West, K. D., Edison, H. J. and Cho, D. (1993). A utility-based comparison of some models of exchange rate volatility. Journal of International Economics, 35(1-2), pp. 23–45.

[47] Zhang, C. and Tu, X. (2016). The effect of global oil price shocks on china's metal markets. Energy Policy, 90, pp. 131–139.