

Do Precious Metal Prices Help in Forecasting South African Inflation?

Mehmet Balcilar^{a,c}, Nico Katzke^{b,*}, Rangan Gupta^c

^a*Department of Economics, Faculty of Business and Economics, Eastern Mediterranean University.*

^b*Department of Economics, Stellenbosch University, South Africa.*

^c*Department of Economics, University of Pretoria, Pretoria, 0002, South Africa.*

Abstract

In this paper we test whether the key metals prices of gold and platinum significantly improve inflation forecasts for the South African economy. We also test whether controlling for conditional correlations in a dynamic setup, using bivariate Bayesian-Dynamic Conditional Correlation (B-DCC) models, improves inflation forecasts. To achieve this we compare out-of-sample forecast estimates of the B-DCC model to Random Walk, Autoregressive and Bayesian VAR models. We find that for both the BVAR and BDCC models, improving point forecasts of the Autoregressive model of inflation remains an elusive exercise. This, we argue, is of less importance relative to the more informative density forecasts. For this we find improved forecasts of inflation for the B-DCC models at all forecasting horizons tested. We thus conclude that including metals price series as inputs to inflation models leads to improved density forecasts, while controlling for the dynamic relationship between the included price series and inflation similarly leads to significantly improved density forecasts.

Keywords:

Bayesian VAR, Dynamic Conditional Correlation, Density forecasting, Random Walk, Autoregressive model

JEL: C11, C15, E17

*Corresponding author

Email addresses: mehmet@mbalcilar.net (Mehmet Balcilar),
nicokatzke@sun.ac.za (Nico Katzke), rangan.gupta@up.ac.za. (Rangan Gupta)

1. Introduction

The value of the local currency in South Africa (SA hereafter) is often linked to a large extent to commodity prices, in particular that of precious metals (which make up nearly a fifth of total exports). As such, markets tend to view precious metal price movements as significant factors explaining domestic currency movements. Currency fluctuations, in turn, impact overall prices in the economy, the extent of which is generally unclear in SA. The purpose of our paper is then to assess whether metals prices should in fact be considered as important inputs in forecasting local inflation.

We consider in this analysis the prices of gold and platinum, which make up the largest part of our precious metals export basket. We set out to test whether the inclusion of these key metal price series improves our ability to forecast inflation for SA. We also test whether explicitly controlling for the time-varying nature of co-movement between these series significantly improves point and density forecasts.

Similar to the work of [Chen, Turnovsky, and Zivot \(2011\)](#), our analysis excludes other fundamental factors which are based on alternative structural models of price dynamics, such as including the output gap or measures of financial development or trade openness.¹ The main objective of this paper is to determine whether gold and platinum prices, which can be considered largely exogenous in terms of local price discovery, are useful in complementing forecasting models of inflation. This will be tested by using out-of-sample point and density forecasts of SA inflation for the period since adopting inflation targeting in 2000. To achieve this, we first fit Bayesian Vector Autoregression (B-VAR) and Bayesian VAR Dynamic Conditional Correlation (B-DCC) models, which we then use to produce out-of-sample forecasts at different horizons. We then compare to the naive Random Walk and Autoregressive model, in terms of the forecasts of inflation, in order to assess whether any meaningful forecasting information has been added by the metals series. This follows as the latter benchmark models are solely based on past information of inflation.

Both the B-VAR and B-DCC models are estimated using Bayesian Markov Chain Monte Carlo (MCMC) methods. Forecasts are generated using recursive estimations, while expanding the estimation sample as forecasting moves

¹C.f. [Ustyugova and Gelos \(2012\)](#) for a structural analysis on the impact of commodity prices on inflation

forward. The Bayesian procedures make use of Normal diffuse priors and posteriors, and the models are estimated using Gibbs sampling.

Our results provide insight into the usefulness of employing Bayesian shrinkage methods to VARs and utilizing time-varying correlation estimates in forecasting inflation using real price inputs, in addition to assessing the importance of precious metals prices in forecasting inflation. The results can then be summarized in two key findings. Firstly, that gold and platinum prices generally provide useful information as input to inflation density forecasts for South Africa at multiple horizons since 2000. This is complementary to findings of [Chen, Turnovsky, and Zivot \(2011\)](#), who also illustrate the importance of considering metals price series as inputs to SA and other inflation targeting emerging market inflation forecasts. Secondly, we find that utilizing time-varying correlation estimates also improves density forecasts of inflation for variables included in the estimation. A future study might consider utilizing similar strategies to test other structural inputs in forecasting inflation.

The paper is organized as follows: Section [2](#) discusses the literature relevant to our study and contextualizes our approach. Thereafter, Section [3](#) discusses the methodology that we will use in order to address the questions posed in the introduction, using the data discussed in [4](#). Section [5](#) discusses the results, after which we conclude the paper in Section [6](#).

2. Literature Review

Most economists and monetary authorities would agree that commodity prices have significant inflationary consequences, although, as suggested by [Gospodinov and Ng \(2013\)](#), opinions on the formal link between inflation and commodity prices remain divided. Some argue that asset market- and commodity prices should be considered leading indicators to the general price level, while others argue that idiosyncratic movements impact prices mainly through the distribution channel. Despite inconclusive evidence of the direct link between commodity prices and inflation², suggestions as to how authorities should respond to commodity price signals remain divided. [Bean \(2004\)](#) provides a more detailed comparison of views on how to approach a build-up of general asset market prices in an inflationary targeting regime. [Fuhrer and](#)

²C.f. [Hooker \(2002\)](#) and [Stock and Watson \(2001\)](#) who suggest that evidence of commodity prices improving inflation forecasts are both elusive and episodic

Moore (1992) and Bernanke and Gertler (2000), e.g., suggest that authorities should not respond to asset market prices as it could lead to a loss in inflationary control.³ Others, such as Cecchetti, Genberg, and Wadhvani (2002), have suggested that policy initiatives aimed at targeting asset price misalignments could improve general price stability and the overall macroeconomic performance.

Despite divergent views on the appropriate actions to be taken by inflation targeting authorities, evidence has been provided as to the importance of certain price indexes in improving general price level forecasts. Gospodinov and Ng (2013) provide evidence that a reduced rank of multiple commodity price indexes, using a principal components approach, produces significant improvements in the predictive power of inflation forecasts. Chen, Turnovsky, and Zivot (2011) consider four commodity exporting emerging markets which have adopted inflation targeting, including SA, and show that considering commodity price aggregates provide predictive power to inflation forecasts. In particular, they highlight the importance of considering metals price series for SA inflationary forecasts.

We build on the work done by Chen et al. (2011), and focus on SA inflation forecasts using two key metals series: gold and platinum. Precious metals typically make up about 6% of SA exports⁴, and as such price fluctuations could be regarded as having a significant impact on currency valuation. This, in turn, might significantly impact the price setting mechanisms in the economy, which we test formally in this paper. We thus test whether two key precious metal prices add to the forecasting power of general price levels in the domestic economy.

Our approach to answering this question differs from Chen et al. (2011) in that we control for the dynamic nature of the co-movements between the price series in our sample. We follow the methodology developed by Della Corte, Sarno, and Tsiakas (2010) in using Bayesian techniques for the estimation of parameters in our DCC model, which is used to estimate time-varying co-dependence structures. Our methodological construct follows that of Lombardi and Ravazzolo (2013), who study the ability of commodity prices in forecasting equity market prices. The authors use bivariate Bayesian VAR

³Bernanke and Gertler (2001) suggests, however, that authorities could respond if such price changes reflect changes in forward inflationary expectations.

⁴According to the Preliminary Statement of Trade Statistics, 2014.

and bivariate Bayesian DCC models to estimate 1,2...24 step ahead point and density forecasts for their studied returns series. They then compare these fits to forecasts from a Random Walk model and Autoregressive model fits. The point and density forecast estimates are then compared using statistical test procedures discussed in Section 3. The authors' findings suggest that the models provide similar point estimates, but that the Bayesian DCC model consistently provides better density forecasts for commodity prices across all the horizons tested. They thus conclude that controlling for time variation in the covariance matrix between the bivariate series pairs, significantly improves density forecasts in their sample.

3. Methodological Discussion

The US Federal Reserve first examined the empirical relationship between commodity price changes and US inflation, using bivariate VAR models (Furlong and Ingenito, 1996). Since then, our understanding of the shortcomings of VAR models has grown, particularly as regards the pitfalls of over-parameterization which could lead to poor results. Several Bayesian type shrinkage techniques have since been introduced, including the Minnesota prior of Doan, Litterman, and Sims (1984), which we will use in calculating our comparative VAR and DCC estimates.⁵

In order to assess whether metals series provide useful information regarding the forecasting of the inflation series for SA, we compare the forecasts of the BVAR and B-DCC models to two benchmark models, the RW and AR models, which have proven in the past to be hard-to-beat in out-of-sample forecasting. As they are both nested in the BVAR model (albeit without the need for Bayesian parameter estimation), we do not discuss these models below.

Our first model that we will use to compare relative to the benchmarks inflation forecasts, is the bivariate Bayesian Vector Autoregressive (BVAR) model. It takes the following form:

$$y_t = c + B(L)y_{t-1} + e_t, \quad e_t \sim N(0, \phi) \quad (1)$$

where y_t is a 2×1 vector for inflation relative to gold and platinum prices, respectively. The errors are also assumed normally distributed with bivariate

⁵c.f. Korobilis (2013) and Baumeister and Kilian (2012) who find that VAR forecasts are significantly improved when using Bayesian shrinkage methods.

covariance matrix ϕ . We find an optimal lag structure of 7 for the inflation model, using the AIC information criterion. This implies that we use for all our BVAR estimates an optimal lag structure of 7, as the BVARs use a prior distribution (discussed below) which matches the in-sample fit of the AR(7) model. The BVAR(7) model is then estimated by setting the priors according to the procedure developed by [Litterman \(1986\)](#), and extended by [Kadiyala and Karlsson \(1997\)](#). This essentially implies using a Minnesota prior, whereby the VAR equations are effectively “centered” around a random walk with drift. As discussed in [Kadiyala and Karlsson \(1997\)](#), this approach effectively shrinks the diagonal elements of B_1 in equation 1 toward unity, and the other parameters to zero. The prior specification also incorporates the belief that autoregressive lags should outweigh other variables’ lags, as well as more recent lags are assumed to outweigh that of earlier lags. The moments for the prior distributions of the coefficients can thus be represented as follows:

$$E[(B_k)_{ij}] = \begin{cases} \delta_i & j = i, k = 1; \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

$$\sigma[(B_k)_{ij}] = \begin{cases} \frac{\lambda^2}{k^2}, & j = 1; \\ \vartheta \cdot \frac{\lambda^2}{k^2} \cdot \frac{\sigma_i^2}{\sigma_j^2}, & \text{otherwise} \end{cases}. \quad (3)$$

Here the coefficients are assumed a priori independently and identically distributed. As the variables are all made stationary by differencing, the prior on the mean is also zero on the first own lag and not unity. We also set δ_i equal to zero to reflect the mean reverting data. Also, the covariance matrix is assumed diagonal and fixed, with $\Sigma = \text{diag}(\sigma_i^2, \sigma_j^2)$, and the intercept prior assumed diffuse. λ can be thought of as how close the prior distribution is to that of the random walk, and thus reflects the importance of prior beliefs to information gleaned from the data. The optimal λ ’s in our estimations were 0.2486 for gold and 0.3959 for the platinum series.⁶ We set it such that the average in-sample fit of the bivariate BVARs, with gold and platinum inputs, have the same fit as the AR(7) model of inflation. $1/k^2$ is the rate at which prior variance decreases as lags decrease, and $\vartheta \in (0, 1)$ governs the relative importance of more recent lags. Similar to [Bańbura et al. \(2010\)](#), which we set ϑ equal to unity to impose a normal inverted Wishart

⁶The higher the value of λ , the closer the posterior would be to the OLS estimates. A smaller value implies the observed data influences the posterior distribution less.

prior. Lastly, the ratio of variances, σ_i^2/σ_j^2 , in equation 3 accounts for scale and variability in the data (our notation follows that of Bańbura, Giannone, and Reichlin (2010) who provide a deeper discussion of the Bayesian VAR approach).

Next we model the bivariate dynamic structure of correlations between inflation and gold and platinum prices, respectively, by using the Dynamic Conditional Correlation (DCC) model extension, developed by Engle (2002), to the B-VAR model described earlier. This model offers a flexible and parsimonious approach to extract time-varying correlations between series. The DCC model extension to the BVAR(7) model describing inflation, y_t , can be written as follows:⁷

$$\begin{aligned}
y_t &= c + B(L)y_{t-1} + e_t \\
e_t &= H_t^{1/2}z_t; \quad z_t \sim N(0, I_N) \quad \& \quad H_t = D_t R_t D_t \\
D_t^2 &= \text{diag}(\sigma_{1,t}, \dots, \sigma_{N,t}) \\
\sigma_{i,t}^2 &= \gamma_i + \kappa_{i,t}v_{i,t-1}^2 + \eta_i\sigma_{i,t-1}^2, \quad \forall i \\
R_{t,i,j} &= \text{diag}(Q_{t,i,j})^{-1}Q_{t,i,j}\text{diag}(Q_{t,i,j})^{-1} \\
Q_{t,i,j} &= (1 - \alpha - \beta)\bar{Q} + \alpha z_t z_t' + \beta Q_{t,i,j}
\end{aligned} \tag{4}$$

H_t is the conditional covariance matrix, and z_t the standard normal disturbances. Also, $\sigma_{i,t}$ is the univariate volatility model for each series i .⁸ $Q_{t,i,j}$ is the unconditional covariance structure and $R_{t,i,j}$ the dynamic conditional correlation estimates between i and j . In our estimation we use only bivariate pairs, thus $N = 2$, and therefore the parameters α and β reduce to scalars. For all our estimations, the following necessary conditions for positive semi-definiteness hold: $\alpha > 0$, $\beta > 0$ and $\alpha + \beta < 1$. Engle (2002) also suggests

⁷Refer to Bauwens et al. (2006) for a detailed discussion on multivariate GARCH model extensions, including the DCC model.

⁸Consistent with most of the literature and to simplify our estimations, we use GARCH(1,1) models for each of the univariate specifications. All parameters are significant and the constraints on the parameters are met. Results are omitted for brevity.

specifying a log-likelihood for the model in 4 defined as:

$$\begin{aligned} \ln(L) = & -\frac{1}{2} \sum_{t=1}^T [N \ln(2\pi) + 2 \ln|D_t| + e_t' D_t^{-1} D_t^{-1} e_t \\ & - \varepsilon_t \varepsilon_t' + \ln|R_{t,i,j}| + \varepsilon_t' R_{t,i,j}^{-1} \varepsilon_t] \end{aligned} \quad (5)$$

We then perform Bayesian estimation of the DCC model parameters to overcome one of the major drawbacks of this approach, its static correlation structure. This is done using the Metropolis Hastings algorithm for estimating the DCC parameters, similar to that suggested by [Della Corte, Sarno, and Tsiakas \(2010\)](#).⁹

After fitting both the BVAR and BVAR-BDCC GARCH models, we proceed to generate point and density inflation forecasts. We compute $h = 1, 2, \dots, 24$ step ahead monthly forecasts for each of the models discussed on an iterative basis, in order to compare it to RW and AR forecasts. The comparisons are made based on the following statistics. Firstly, we compare the point forecast estimates using the familiar Root Mean Squared Error (RMSE) statistic, which is given as:

$$RMSE_{k,h} = \sqrt{\sum_{t=\underline{t}}^{\bar{t}-h} \frac{(y_{t+h,h} - \tilde{y}_{t+h,k})^2}{t^*}} \quad (6)$$

where $t^* = \bar{t} - h + \underline{t} + 1$, \underline{t} is the beginning and \bar{t} the end of the forecasting period, and \tilde{y}_t and y_t the forecast and true values, respectively. Using this measure, we can assess which model provides the best point forecast estimate.

We also compare density forecasts using the Logarithmic Score statistic, LS, as discussed in detail in [Mitchell and Hall \(2005, 2007\)](#). The usefulness of evaluating density forecasts in addition to point estimates has been outlined by [Tay and Wallis \(2000\)](#), who argue that density measures allows for a full impression of the uncertainty associated with forecasts. In our analysis, it provides an estimate of the probability distribution of future values for inflation as predicted by each model. Following the discussion in [Mitchell and Hall](#)

⁹For a detailed discussion of the Bayesian DCC approach, the reader is referred to [Della Corte et al. \(2010\)](#), while [Lombardi and Ravazzolo \(2013\)](#) provides a good concise overview of this approach in their appendix.

(2005), density forecasts can be combined using “optimal” weights. These optimal weights minimize the distance between candidate model k ’s combined density forecasts, $p(\tilde{y}_{k,T+1}|y_{1:t})$, and the true densities, $p(y_{k,T+1}|y_{1:t})$, which are unknown. In doing so, the Kullback-Leibler Information Criterion (KLIC) measure is used to obtain these optimal weights for model k , which minimize the distance between the combined and true densities (Mitchell and Hall, 2007, p.4):

$$\overline{KLIC}_{k,t+h} = \int p(y_{t+h}|y_{1:t}) \ln \frac{p(y_{t+h}|y_{1:t})}{p(\tilde{y}_{t+h}|y_{1:t})} \quad (7)$$

$$\overline{KLIC}_{k,t+h} = E_t \{ (\ln p(y_{t+h}|y_{1:t}) - \ln p(\tilde{y}_{t+h}|y_{1:t})) \} \quad (8)$$

where $E_t = E(.|F_t)$ is the conditional expectation given information set F_t . The smaller the distance in 8, the closer the density forecast of model k is to the “true density”. It then follows that under some regularity conditions, the criterion can be rewritten as¹⁰:

$$\overline{KLIC}_{k,t+h} = \frac{1}{t^*} \sum_{t=\underline{t}}^{\bar{t}-h} (\ln p(y_{t+h}|y_{1:t}) - \ln p(\tilde{y}_{t+h}|y_{1:t})) \quad (9)$$

where t^* corresponds to our out-of-sample range discussed in Section 5, and the sample statistic, $\frac{1}{t^*} \sum(.)$, is used as an unbiased estimate of E_t . Intuitively, the KLIC statistic thus chooses the model which, on average, gives the highest probability to the events which occurred (or the model having the highest posterior probability). Minimizing the KLIC statistic, is equivalent to maximizing the second part of equation 9, which is the Logarithmic Score statistic. This follows as the “true” density is not observable, yet we need not know $\ln p(y_{t+h}|y_{1:t})$ in order to compare two models’ density forecast fits, as it would be relative to the same “true” density. Thus, we have:

$$LS_k = -\frac{1}{t^*} \sum_{t=\underline{t}}^{\bar{t}-h} \ln p(\tilde{y}_{t+h}|y_{1:t}) \quad (10)$$

In order to assess the significance of the model performance in our out-of-sample point forecast estimates, we use Diebold and Mariano (1995)’s (DM)

¹⁰C.f. Mitchell and Hall (2007, p. 4). Here we follow the notation of Lombardi and Ravazzolo (2012)

test. The test evaluates the null hypothesis of equal forecasting accuracy of each bivariate model pair compared to the alternative that one of the models outperformed the other in terms of forecasting accuracy. We use the small sample adjusted version of the DM test, as suggested by [Harvey, Leybourne, and Newbold \(1997\)](#), which is a pairwise test that adjusts the DM, and is defined as follows:

$$E[d_{i,t}] = E[\Lambda_{1,t}^w - \Lambda_{2,t}^w] = 0 \quad (11)$$

and then following [Harvey et al. \(1997\)](#), we adjust it as follows:

$$\text{Adjusted } DM = \left(\frac{t^* + 1 - 2h + t^{*-1}h(h-1)}{t^*} \right)^{1/2} \hat{V}(\bar{d}_i^{-1/2})\bar{d}_i \quad (12)$$

where $\Lambda_{i,t}^w = y_t - \tilde{y}_{i,t}$ is the weighted loss function as described in [Dijk and Franses \(2003\)](#), h being the forecast horizon, and $\hat{V}(\bar{d}_i^{-1/2})$ the estimated variance of series $d_{i,t}$. The Adjusted DM test statistic is then compared to critical values from the t -distribution with degrees of freedom $t^* - 1$. In the adjusted DM above, we use Newey-West HAC consistent variance estimators. The Bandwidth for the Newey-West estimator is then selected using Andrews' (1991) automatic bandwidth selector.¹¹

In order to compare model significance in terms of density forecasts, we follow a similar yet somewhat adjusted approach to the above, as outlined in [Mitchell and Hall \(2005\)](#). The null hypothesis of equal density forecast is:

$$E[d_{i,t}] = 0 \Rightarrow KLIC = 0 \quad (13)$$

with the sample mean \bar{d} defined as ([Mitchell and Hall, 2005](#), p.1004):

$$\bar{d} = KLIC = \frac{1}{t^*} \sum_{t=\bar{t}}^{\bar{t}-h} [\ln p_t(z_{1t})^* - \ln \phi(z_{1t}^*)] \quad (14)$$

with $z_{1t} = \int_{-\infty}^{y_t} g_u(u)du$, and $z_{1t}^* = \Phi^{-1}z_{1t}$, where $\phi(\cdot)$ is the standard normal density function and Φ the c.d.f of the standard normal. As noted in [Mitchell](#)

¹¹For details, see: Andrews, D. W. (1991). Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica: Journal of the Econometric Society*, 817-858.

and Hall (2005), testing for the departure of $\{z_{1t}\}_{t=1}^{t^*}$ from i.i.d $N(0, 1)$ ¹² is equivalent to testing the departure of the forecasted density from the true density, $p(y_{t+h}|y_{1:t})$, as defined earlier.

In order to test equation 14, we follow Mitchell and Hall (2005) in using the central limit theorem and, under appropriate assumptions, testing the distribution:

$$\sqrt{t^*}(\bar{d} - E(d_t)) \rightarrow N(0, \Omega) \quad (15)$$

where the general representation for the covariance matrix, Ω , is given in West (1996), and reduces to the long run covariance matrix, S_d . This is, in turn, associated with d_t in that it is 2π times the spectral density of $d_t - E(d_t)$ at zero frequency (Mitchell and Hall, 2005).¹³ This test then reduces to the small sample corrected and HAC robust DM test if d_t tested for point error forecasts. As argued in Mitchell and Hall (2005), such testing for the significance of departures of z_{1t}^* from $N(0, 1)$, is both easier and more sensible than testing the distance of density forecasts as defined in equation 9.

The results of both the point and density forecast tests are reported in Table A.4 and discussed in Section 5.

4. Data

In this study we use monthly data for the period 1968M2 to 2013M12, for SA inflation (CPI) as well as gold and platinum price indexes. Our data for CPI inflation was obtained from the Global Financial database, while the platinum and gold price series were obtained from www.kitco.com. We used the dollar prices for both gold and platinum, firstly as these metals are traded globally in dollars, and secondly in order to avoid exchange rate effects on inflation in our estimates. Table A.1 shows the descriptive statistics of the returns to each of the series¹⁴, and figure B.1 graphs the month-on-month inflation and continuously compounded returns for the gold- and platinum series. From both Table A.1 and figure B.1 it is clear that platinum prices

¹²Which is, in turn, equivalent to testing z_{1t} 's departure from i.i.d. $U(0, 1)$

¹³More details of estimating S_d follows in Mitchell and Hall (2005, p. 1004 - 1005).

¹⁴For each price index, we use the first difference of the logarithmic transformation to achieve stationarity.

show the greatest variability, while both metals series show far more price variability than prices on aggregate.

Also, all the price returns series reject the null of normality, and display excess kurtosis and skewness. Finally, all three series show logarithmically differenced stationarity, as can be seen from the Augmented Dickey-Fuller (ADF) and Philips-Perron (PP) tests contained in the appendix (see Table A.2). Although less convincing for inflation using the standard ADF test, the PP test indeed shows clear evidence of stationarity.

Figure B.2 in the appendix compares the bivariate Bayesian Dynamic Conditional Correlations (B-DCC) between inflation and the respective metals price series, relative to their unconditional sample correlation. From the figures it is clear that the unconditional estimates for both are consistently lower than for the dynamic estimates. We also note from these figures that the correlation between the inflation series and the metals returns series typically display relatively muted and positive correlations. As such, we require the point and density forecast evaluation results to be able to infer whether metals price series should indeed be considered as important inputs to inflation forecasts.

5. Results

Table A.3 shows both the point and density out-of-sample forecast accuracy tests for the B-VAR and BVAR B-DCC models discussed in Section 3. The period considered ranges from 2000M01 – 2013m12, which corresponds to the period of inflation targeting.

From Table A.3, it is clear that the results for the point forecasts suggest that the AR(7) model is particularly difficult to improve upon. This is consistent with the findings of Lombardi and Ravazzolo (2012) and most other studies that similarly have difficulty in finding more accurate point estimates than the autoregressive models. Table A.4 confirms the significance of the out performance relative to both BDCC estimates, using gold and platinum as bivariate inputs, in terms of point forecasts.

Nonetheless, as argued in Section 3, we should be more interested in the ability of models to provide accurate density forecast estimates. This follows as it is a more holistic approach to assessing the model’s ability to predict future price movements. The LS statistic in Table A.3, in contrast to the findings for point forecasts, shows that the B-DCC models consistently provide improved density forecasts of inflation at every horizon versus all the

other models. Also, the B-VAR for both metals series also provide improved density forecasts versus the RW and AR(7) benchmarks. Collectively, this implies that at all horizons we see that controlling for dynamic correlations, using BDCC extensions to the BVAR model, improves density forecasts. We also deduce that metals price series contain information which improves inflation forecasting power at all horizons as well. The improvement from the BVAR as a result of the addition of BDCC estimates, are more significant as the forward horizon increases, as can be seen from Table A.4. For all density forecast horizons we see that the BDCC models provide significantly more accurate estimates than the RW model, while for most AR(7) horizons the tests suggest significant improvements at longer horizons.

Our results can then be summarized as follows. Firstly, metals price series, proxied for by gold and platinum prices, provide useful information in order to improve density forecasts of inflation based solely on past inflation information. Secondly, improving point forecast estimates from naive models remain an elusive, albeit less informative, exercise. Lastly, we see that controlling for the dynamic relationship between inflation and other price series used as model inputs, leads to statistically significant improvements in density forecasts.

6. Conclusion

This paper studies the dynamic relationship between inflation and metals price series for South Africa, for the period 1968 – 2013. It tests two hypotheses. First, whether metals series provide useful information in the point and density out-of-sample forecasting of inflation for the period since adopting inflation targeting. Secondly, we test whether controlling for the dynamic conditional co-movements between the bivariate pairs of price series yields improved forecasting performance. This is done using a bivariate Bayesian VAR Bayesian DCC-GARCH model, and comparing the out-of-sample forecasting performances at various horizons. To achieve this, we make use of the RMSE statistics for point forecast estimates, the Log-Score statistic for density forecasts and a modified version of the Diebold and Mariano (1995) test and the Mitchell and Hall (2005) test to evaluate the statistical significance of model forecast out-performance for both point and density forecasts, respectively.

Our findings can be summarized as follows. We find that improving point forecasts of naive Random Walk and Autoregressive models for the inflation series, using our dynamic model estimates, remain an elusive exercise. This, we argue, is of less importance relative to the more informative density forecasts, for which we find significantly improved forecasts of inflation for both the BVAR and BDCC models over the naive benchmarks. Particularly, we find that controlling for dynamic correlations between the series leads to superior density forecasts at all horizons. This allows us to make several conclusions. Firstly, that including metals price series as inputs to inflation models leads to improved density forecasts. Secondly, controlling for the dynamic relationship between the included price series and inflation similarly leads to significantly improved forecasts. This implies that forecasters should consider the importance of metals series in describing the movements of prices in South Africa, while advanced models studying future price movements should ideally control for the dynamic relationships of the series considered. Our Bayesian estimates allow the dynamic structure of correlation to be more flexible and adjustable relative to the standard DCC estimates typically used in the literature.

References

- Bañbura, M., D. Giannone, and L. Reichlin (2010). Large bayesian vector auto regressions. *Journal of Applied Econometrics* 25(1), 71–92.
- Baumeister, C. and L. Kilian (2012). Real-time forecasts of the real price of oil. *Journal of Business & Economic Statistics* 30(2), 326–336.
- Bauwens, L., S. Laurent, and J. V. Rombouts (2006). Multivariate garch models: a survey. *Journal of applied econometrics* 21(1), 79–109.
- Bean, C. (2004). Inflation targeting: the uk experience. *Perspektiven der Wirtschaftspolitik* 5(4), 405–421.
- Bernanke, B. and M. Gertler (2000). Monetary policy and asset price volatility. Technical report, National bureau of economic research.
- Bernanke, B. S. and M. Gertler (2001). Should central banks respond to movements in asset prices? *american economic review*, 253–257.

- Cecchetti, S. G., H. Genberg, and S. Wadhvani (2002). Asset prices in a flexible inflation targeting framework. Technical report, National Bureau of Economic Research.
- Chen, Y.-c., S. J. Turnovsky, and E. Zivot (2011). Forecasting inflation using commodity price aggregates. Technical report, University of Washington, Department of Economics.
- Della Corte, P., L. Sarno, and I. Tsiakas (2010). Correlation timing in asset allocation with bayesian learning. Technical report, Warwick Business School,.
- Diebold, F. X. and R. S. Mariano (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics* 13(3).
- Dijk, D. and P. H. Franses (2003). Selecting a nonlinear time series model using weighted tests of equal forecast accuracy*. *Oxford Bulletin of Economics and Statistics* 65(s1), 727–744.
- Doan, T., R. Litterman, and C. Sims (1984). Forecasting and conditional projection using realistic prior distributions. *Econometric reviews* 3(1), 1–100.
- Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics* 20(3), 339–350.
- Fuhrer, J. and G. Moore (1992). Monetary policy rules and the indicator properties of asset prices. *Journal of Monetary Economics* 29(2), 303–336.
- Furlong, F. and R. Ingenito (1996). Commodity prices and inflation. *Economic Review-Federal Reserve Bank of San Francisco*, 27–47.
- Gospodinov, N. and S. Ng (2013). Commodity prices, convenience yields, and inflation. *The Review of Economics and Statistics* 95(1), 206–219.
- Harvey, D., S. Leybourne, and P. Newbold (1997). Testing the equality of prediction mean squared errors. *International Journal of forecasting* 13(2), 281–291.

- Hooker, M. A. (2002). Are oil shocks inflationary?: Asymmetric and nonlinear specifications versus changes in regime. *Journal of Money, Credit, and Banking* 34(2), 540–561.
- Kadiyala, K. and S. Karlsson (1997). Numerical methods for estimation and inference in bayesian var-models. *Journal of Applied Econometrics* 12(2), 99–132.
- Korobilis, D. (2013). Var forecasting using bayesian variable selection. *Journal of Applied Econometrics* 28(2), 204–230.
- Litterman, R. B. (1986). Forecasting with bayesian vector autoregressions five years of experience. *Journal of Business & Economic Statistics* 4(1), 25–38.
- Lombardi, M. J. and F. Ravazzolo (2012). Oil price density forecasts: exploring the linkages with stock markets.
- Lombardi, M. J. and F. Ravazzolo (2013). On the correlation between commodity and equity returns: implications for portfolio allocation. *BIS Working Paper*.
- Mitchell, J. and S. G. Hall (2005). Evaluating, comparing and combining density forecasts using the klic with an application to the bank of england and niesr fancharts of inflation*. *Oxford bulletin of economics and statistics* 67(s1), 995–1033.
- Mitchell, J. and S. G. Hall (2007). Combining density forecasts. *International Journal of Forecasting* 23(1), 1–13.
- Stock, J. H. and M. W. Watson (2001). Forecasting output and inflation: the role of asset prices. Technical report, National Bureau of Economic Research.
- Tay, A. S. and K. F. Wallis (2000). Density forecasting: a survey. *Journal of forecasting* 19(4), 235–254.
- Ustyugova, Y. and G. Gelos (2012). Inflation responses to commodity price shocks-how and why do countries differ? *IMF Working Paper*.
- West, K. D. (1996). Asymptotic inference about predictive ability. *Econometrica: Journal of the Econometric Society*, 1067–1084.

Appendix A. Tables

Table A.1: Descriptive Statistics of the returns series

	Inflation	Gold	Platinum
Mean	0.739599	0.644396	0.33376
Median	0.654025	0.126296	0.20334
Maximum	4.211154	39.46875	28.47158
Minimum	-0.74219	-18.3862	-29.2945
Std. Dev.	0.637043	4.924338	6.301847
Skewness	1.107125	1.256602	0.28834
Kurtosis	5.768518	11.20655	6.555835
Jarque-Bera	288.5311	1691.196	297.9193
Probability	0.000	0.000	0.000
Observations	551	551	551

Table A.2: Tests for Unit Roots

Inflation	t-Statistic	Prob.
Augmented Dickey-Fuller test statistic	-2.774	0.0628
Test critical values: 1% level	-3.442	
5% level	-2.867	
10% level	-2.57	
Gold Returns	t-Statistic	Prob.
Augmented Dickey-Fuller test statistic	-15.97	0.000
Test critical values: 1% level	-3.442	
5% level	-2.867	
10% level	-2.57	
Platinum Returns	t-Statistic	Prob.
Augmented Dickey-Fuller test statistic	-18.26	0.000
Test critical values: 1% level	-3.442	
5% level	-2.867	
10% level	-2.57	
Inflation	Adj. t-Statistic	Prob.
Phillips-Perron test statistic	-22.89	0.000
Test critical values: 1% level	-3.442	
5% level	-2.867	
10% level	-2.57	
Gold Returns	Adj. t-Statistic	Prob.
Phillips-Perron test statistic	-17.51	0.000
Test critical values: 1% level	-3.442	
5% level	-2.867	
10% level	-2.57	
Platinum Returns	Adj. t-Statistic	Prob.
Phillips-Perron test statistic	-18.22	0.000
Test critical values: 1% level	-3.442	
5% level	-2.867	
10% level	-2.57	

Both the Augmented Dickey-Fuller and Phillips-Perron tests use MacKinnon's (1996) one-sided p-values. The latter test also employs the Newey-West automatic bandwidth selector and Bartlett's kernel.

Table A.3: Forecast Comparisons using RMSE and Logarithmic Score Tests

h=1¹						
	RW	AR	B-DCC Gold	B-DCC Plat	BVAR Gold	BVAR Plat
RMSE ²	0.5356	0.4722	0.4875	0.4719	0.4804	0.4667
LS ³	-0.9524	-0.7355	-0.7155	-0.71	-0.7293	-0.7128
h=3						
	RW	AR	B-DCC Gold	B-DCC Plat	BVAR Gold	BVAR Plat
RMSE	0.5940	0.4869	0.5004	0.5038	0.4938	0.4886
LS	-1.3876	-0.7619	-0.7417	-0.7471	-0.7558	-0.7508
h=6						
	RW	AR	B-DCC Gold	B-DCC Plat	BVAR Gold	BVAR Plat
RMSE	0.5721	0.5091	0.5231	0.5211	0.5135	0.5105
LS	-1.6938	-0.7982	-0.7641	-0.7705	-0.7871	-0.7857
h=12						
	RW	AR	B-DCC Gold	B-DCC Plat	BVAR Gold	BVAR Plat
RMSE	0.5605	0.5407	0.5542	0.5548	0.5419	0.5427
LS	-2.0181	-0.8344	-0.7678	-0.7666	-0.8177	-0.8186
h=24						
	RW	AR	B-DCC Gold	B-DCC Plat	BVAR Gold	BVAR Plat
RMSE	0.5551	0.5232	0.5427	0.5421	0.5230	0.5235
LS	-2.3539	-0.8791	-0.8171	-0.8087	-0.8754	-0.8749

¹ h indicates the forecast horizon. For brevity we include only some of the forecast horizons, with the other estimates available from the authors upon request, although they do not change the results in any way.

² RMSE is the Root Mean Square Error for point forecasts. A lower statistic indicates a better point forecast.

³ LS is the log score statistic for density forecasts. A higher LS statistic means better density coverage.

Table A.4: Point and Density Forecast Accuracy

Gold point forecasts ¹			
horizon	RW vs. BDCC	AR vs. BDCC	BDCC vs. BVAR
1	1.149	-2.429**	2.139**
3	2.175**	-1.584	1.524
6	1.291	-1.797*	2.567**
12	0.135	-3.544***	3.313***
24	0.214	-4.216***	4.249***
Gold density forecasts ²			
horizon	RW vs. BDCC	AR vs. BDCC	BDCC vs. BVAR
1	-3.949***	-0.232	0.061
3	-9.286***	-0.46	0.328
6	-10.999***	-1.141	1.137
12	-12.844***	-2.883***	2.901***
24	-35.826***	-1.851*	1.604
Platinum point forecasts			
horizon	RW vs. BDCC	AR vs. BDCC	BDCC vs. BVAR
1	1.611	0.052	1.56
3	2.205**	-3.208***	3.747***
6	1.343	-2.104**	2.146**
12	0.12	-3.493***	3.427***
24	0.222	-2.880***	2.916***
Platinum density forecasts			
horizon	RW vs. BDCC	AR vs. BDCC	BDCC vs. BVAR
1	-4.723***	-0.605	-0.271
3	-9.703***	-0.448	0.032
6	-11.122***	-1.216	0.904
12	-12.480***	-2.789***	2.942***
24	-35.687***	-1.942*	1.696*

¹ The adjusted Diebold-Mariano (DM) test is used for the point forecast evaluation. A **positive** and significant value implies model 2 provides a significantly better forecast, and *vice versa*.

² The [Mitchell and Hall \(2005\)](#) test is used for density forecast evaluation. A **negative** and significant value implies model 2 provides significantly better density forecast.

*, **, and *** denote significance at 10%, 5% and 1%, respectively.

Appendix B. Figures

Figure B.1: Month-on-month Inflation, Gold and Platinum returns




Figure B.2: Bivariate Bayesian DCC estimates

