Resources Policy 45 (2015) 299–306

Contents lists available at ScienceDirect

Resources Policy

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

A real-time quantile-regression approach to forecasting gold returns under asymmetric loss

Christian Pierdzioch*, Marian Risse, Sebastian Rohloff

Department of Economics, Helmut Schmidt University, Holstenhofweg 85, P.O.B. 700822, 22008 Hamburg, Germany

ARTICLE INFO

Article history:
Received 17 March 2015
Received in revised form
30 May 2015
Accepted 13 July 2015
Available online 28 July 2015

JEL classification:
C53
E44

Keywords:
Quantile regression
Forecasting
Asymmetric loss
Gold returns

ABSTRACT

We propose a real-time quantile-regression approach to analyze whether widely studied macroeconomic and financial variables help to forecast out-of-sample gold returns. The real-time quantile-regression approach accounts for model uncertainty, model instability, and the possibility that a forecaster has an asymmetric loss function. Forecasts are computed and evaluated using the same asymmetric loss function. When the loss function implies that an underestimation is somewhat more costly than an overestimation of the same size, the forecasts computed using the real-time quantile-regression approach outperform forecasts implied by an autoregressive benchmark model.

1. Introduction

Against the background of recent financial market turbulences, research on the determinants of gold returns has mushroomed. Among the determinants that researchers have studied are the inflation rate (Beckmann and Czudaj, 2013a; Batten et al., 2014), the oil price (Zhang and Wei, 2010; Reboredo, 2013b), the exchange rate (Pukthuanthong and Roll, 2011; Reboredo, 2013b), and business-cycle fluctuations (Pierdzioch et al., 2014b). Studying whether the various determinants studied in earlier literature help to forecast gold returns is important because the properties of gold as a safe-haven investment, a low-correlation portfolio diversifier, and a hedge against fluctuating currency values have received much attention in recent research (Hillier et al., 2006; Joy, 2011; Ciner et al., 2013, to name just a few).

Despite many research efforts, no consensus has emerged regarding the core determinants of gold returns. As a result, researchers have studied gold returns by applying flexible forecasting approaches that account for model uncertainty and model instability (see also Vrugt et al., 2007; Aye et al., 2015; Baur et al., 2014; Pierdzioch et al., 2014a, 2014b, 2015). Model uncertainty arises because gold returns may be linked to a potentially large number of determinants, none of which can be excluded a priori on economic grounds. Model instability arises because the relative importance of these determinants most likely has changed over time (Baur, 2011; Batten et al., 2014) and may be state dependent (Wang and Lee, 2011; Wang et al., 2011).

We contribute to earlier literature in that we propose a real-time quantile-regression approach to forecast gold returns. A quantile-regression approach renders it possible to compute forecasts that target the conditional quantiles rather than the conditional mean of the distribution of gold returns. Quantile regressions have received growing attention in the recent finance literature (Basset and Chen, 2001; Engle and Manganelli, 2004; Chuang et al., 2009; Baur et al., 2012) and have been studied recently in a forecasting context by Meligkotsidou et al. (2014), Manzan (2015), and Pedersen (2015). Quantile regressions also have been applied to study gold returns. Ma and Patterson (2013) apply quantile regressions to study the links between the gold price and its macroeconomic and financial determinants. Mensi et al. (2014) use quantile regressions to study how emerging-market stock-market returns depend on gold returns and other macroeconomic and financial factors. Dee et al. (2013) use quantile regressions to explore the link between gold returns, stock-market movements and inflation, and Zagaglia and Marzo (2013) use quantile regressions to study the link between gold returns and exchange-rate movements. Jeong et al. (2012) develop a test for Granger causality in conditional quantiles and apply their test to study the causal links between gold returns, oil-price returns, and

* Corresponding author.
E-mail address: c.pierdzioch@hsu-hh.de (C. Pierdzioch).

http://dx.doi.org/10.1016/j.resourpol.2015.07.002
0301-4207/

exchange-rate movements. Baur (2013) uses quantile regressions to study the link between gold excess returns and the excess returns on a commodity index. Ciner (2015) shows that the quantile regressions render it possible to recover links between, on one hand, CAPM betas and returns of stocks of precious metal mining firms and, on the other hand, trading volume.

Forecasting the conditional quantiles of the distribution of gold returns is a natural forecasting strategy if a forecaster has an asymmetric loss function (for an illustration, see Koenker and Hallock, 2001, p. 146). An asymmetric loss function easily arises in a risk-management context, or simply because of behavioral biases or strategic behavior of forecasters (Laster et al., 1999; Pierdzioch et al., 2013). In recent research on the determinants of gold returns, Pierdzioch et al. (2014b) use asymmetric loss functions to measure the accuracy of out-of-sample forecasts of gold returns. Building on research by Campbell and Thompson (2008), they evaluate forecasts using an out-of-sample R^2 statistic that can be computed under a symmetric and an asymmetric loss function. The computation of forecasts, however, uses the real-time forecasting approach developed by Pesaran and Timmermann (1995, 2000). This approach accounts for model uncertainty and model instability but rests on the assumption that a forecaster has a symmetric loss function because forecasting regressions are estimated by the ordinary-least-squares technique. Hence, in case a forecaster has an asymmetric loss function, the problem arises that the loss function used to compute forecasts differs from the loss function used to evaluate forecasts. The real-time quantile-regression approach that we study in this research overcomes this problem because, as Koenker and Machado (1999) have shown, the potentially asymmetric loss function used for forecast computation can also be used for forecast evaluation. The real-time quantile-regression approach, thus, is an integrated approach to forecasting and evaluating gold returns under asymmetric loss.

We organize the remainder of this research as follows. In Section 2, we outline the real-time quantile-regression approach and we describe how we evaluate forecasts under an asymmetric loss function. In Section 3, we describe our data and we lay out our empirical results. In Section 4, we conclude.

2. The real-time quantile-regression approach

We assume that a forecaster considers n macroeconomic and financial variables, $x_{j,t}$, $j = 1, \ldots, n$, as potential predictors for gold returns, r_{t+1}, in period of time $t+1$. The forecasting model is of the general format $r_{t+1} = \beta_0 + \beta_1 x_{1,t} + \cdots \beta_n x_{n,t} + u_{t+1}$, where u_{t+1} = disturbance term and β_j, $j = 0, 1, 2, \ldots, n$ are regression coefficients to be estimated.

A key problem is that a forecasting model that features all predictor variables is not necessarily the best forecasting model. In principle, a forecaster can choose, in every period of time, t, among the competing forecasting models that feature alternative combinations of the predictor variables (Pesaran and Timmermann, 1995, 2000). Accordingly, we account for model uncertainty by estimating in every period of time, t, all possible combinations of forecasting models, given the n predictor variables.

We use a quantile-regression approach (Koenker and Basset, 1978, for a textbook exposition, see Koenker, 2005) to estimate the forecasting models and, thereby, take into account that a forecaster may have an asymmetric loss function. The following period-loss function forms the foundation of the quantile-regression approach

$$\mathcal{L}(\alpha, \hat{u}_{jt+1,m,\alpha}) = \hat{u}_{t+1,m,\alpha}(\alpha - \mathbf{1}(\hat{u}_{t+1,m,\alpha} < 0)), \tag{1}$$

where $\mathbf{1}(\cdot)$ = indicator function, and $\hat{u}_{t,m,\alpha}$ = forecast error for model m in period of time t, given the quantile parameter, $\alpha \in (0, 1)$. The forecast error is defined as actual returns minus the forecast. If $\alpha=0.5$, the loss function is symmetric in the absolute forecast error, while for $\alpha < 0.5$ ($\alpha > 0.5$) the loss of a negative (positive) forecast error exceeds the loss of a positive (negative) forecast error. In the symmetric case with $\alpha=0.5$, a forecaster should target the median of the distribution of gold returns. If the quantile parameter assumes a value $\alpha < 0.5$ ($\alpha > 0.5$), a forecaster should target the α-quantile of the distribution of gold returns, requiring a downward (upward) adjustment of forecasts to make positive (negative) forecast errors more likely than in the case of $\alpha=0.5$.

Given a quantile parameter, α, we sum up over the period-loss functions to compute the total loss and choose, for every model, m, the parameters, β_α, to minimize

$$\mathcal{L}(\alpha, m, t) = \min_{\beta_\alpha} \sum_{j=01}^{t} \mathcal{L}(\alpha, \hat{u}_{j+1,m,\alpha}), \tag{2}$$

where t denotes the latest period of time for which data that can be used to forecast gold returns are available, and $\hat{u}_{j+1,m,\alpha}$ is interpreted as the in-sample forecast error. The notation β_α emphasizes that the parameters of the forecasting models can differ across quantiles.

In every period of time, t, we select an optimal forecasting model by comparing the m estimated models with a benchmark model, b. To this end, we compute $\min c_{\alpha,m,t}$, with $c_{\alpha,m,t} = \gamma_{m,\alpha}\mathcal{L}(\alpha, m, t)/\mathcal{L}^b(\alpha, t)$, where $\gamma_{m,\alpha} = (t-1)/(t - l_{\beta,m,\alpha})$ penalizes model complexity, $l_{\beta,m,\alpha}$ = length of the vector of regression parameters (neglecting the constant) for model m given the quantile parameter, α, and $\mathcal{L}^b$ = loss under a benchmark model (autoregressive model of order one). In addition, we study various forecast-averaging schemes. Specifically, we compute the mean and the median of the out-of-sample forecasts implied by the m estimated models, and we compute a weighted out-of-sample forecast using, for simplicity, $1/c_{\alpha,m,t}$ as weights (weights are scaled to sum up to unity across models; for other averaging schemes, see Meligkotsidou et al., 2014).

A quantile regression captures potential shifts in the links between gold returns and its macroeconomic and financial determinants across the distribution of gold returns. However, the links between gold returns and its determinants may also change over time due to, for example, financial crises and structural breaks. We account for the resulting model instability by recursively reestimating all possible combinations of forecasting models in every period of time, t, as new data become available. Nicolau and Palomba (2015) argue that using a recursive rather than a rolling-window estimation approach has the advantages that there is no need to specify the length of the rolling window and the information used for estimation is maximized since no observations are dropped to fix the length of the rolling window.

For forecast evaluation, we use the out-of-sample R^2 statistic studied in the context of forecasting gold returns by Pierdzioch et al. (2014b). Their out-of-sample R^2 statistic is similar to the goodness-of-fit criterion for quantile regressions proposed by Koenker and Machado (1999), and it extends the out-of-sample R^2 statistic analyzed by Campbell and Thompson (2008) to the case of an asymmetric loss function. For our quantile-regression approach, the out-of-sample R^2 statistic is given by $R^2(\alpha, b) = 1 - \mathcal{L}(\alpha)/\mathcal{L}^b(\alpha)$, where $\mathcal{L}(\alpha)$ = sum of the out-of-sample losses, and $\mathcal{L}^b(\alpha)$ = sum of the out-of-sample losses for a benchmark model. Hence, the loss function used for forecast evaluation is identical to the loss function used for forecast computation. Given an α-quantile, the cumulated loss, $\mathcal{L}(\alpha)$, is computed as the sum of losses implied by the one-period-ahead forecast errors obtained either from a