<https://helda.helsinki.fi>

Noncausality and the Commodity Currency Hypothesis

Lof, Matthijs Harm

2017-06

Lof , M H & Nyberg , H K 2017 , ' Noncausality and the Commodity Currency Hypothesis ' ,
Energy Economics , vol. 65 , pp. 424-433 . <https://doi.org/10.1016/j.eneco.2017.05.024>

<http://hdl.handle.net/10138/302628>

<https://doi.org/10.1016/j.eneco.2017.05.024>

acceptedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

Noncausality and the Commodity Currency Hypothesis *

Matthijs Lof[†] Henri Nyberg[‡]

Abstract

This paper provides new evidence on the role of exchange rates in forecasting commodity prices. Consistent with previous studies, we find that commodity currencies hold out-of-sample predictive power for commodity prices when using standard linear predictive regressions. After we reconsider the evidence using noncausal autoregressions, which provide a better fit to the data and are able to accommodate the effects of nonlinearities and omitted variables, the predictive power of exchange rates disappears.

Keywords: Commodity prices, exchange rates, noncausal autoregression, nonlinearity

JEL classification: C53, F37, Q02

*We thank Richard Tol (the Editor), two anonymous referees, Helinä Laakkonen, Markku Lanne, Mikko Niemenmaa, Andrew Patton, Barbara Rossi and conference and seminar participants at Aalto University, the Financial Econometrics and Empirical Asset Pricing Conference (Lancaster, 2016), CREST Paris (2015), the 9th and 10th International Conference on Computational and Financial Econometrics (London, 2015; Seville, 2016), the 10th Energy and Finance Conference (London, 2015), the Society for Nonlinear Dynamics and Econometrics (Oslo, 2015), the Nordic Econometric Meeting (Helsinki, 2015) and the Finnish Economic Society (Helsinki, 2015) for useful discussions. The second author is grateful for financial support from the Academy of Finland and the Research Funds of the University of Helsinki.

[†]*Corresponding author.* Aalto University School of Business. E-mail: matthijs.lof@aalto.fi, P.O.Box 21220, 00076 Aalto, Finland, tel:+358504014448.

[‡]University of Turku (Department of Mathematics and Statistics) and University of Helsinki (Department of Political and Economic Studies and Helsinki Center of Economic Research). E-mail: henri.nyberg@utu.fi.

1 Introduction

Understanding the dynamics of commodity prices is of interest not only to commodity traders, but also to policy makers in both commodity exporting and importing countries. Unfortunately, prices of commodities are notoriously hard to forecast. As for many other assets traded on competitive financial markets, commodity prices display near-random walk behavior with changes being virtually unpredictable by past prices or other observable factors¹.

Against this background, Chen, Rogoff and Rossi (2010 - hereafter CRR) find a surprising novel channel of commodity price predictability. 'Commodity currencies', the exchange rates of commodity exporters, appear relevant predictors for commodity prices at quarterly horizons. Using linear predictive regressions, CRR document predictive power both in and out of sample. In an update of the original article, CRR (2014) show this predictability to hold as well for an extended sample period including the recent financial crisis.

Following Engel and West (2005), CRR arrive at the hypothesis that exchange rates predict commodity prices from a standard present value model for exchange rates in which exchange rates (s_t) represent discounted expected fundamentals (f_t):

$$s_t = \gamma \sum_{i=0}^{\infty} \delta^i E_t(f_{t+i}). \quad (1)$$

Campbell and Shiller (1987) show that this present value relation implies that future fundamentals are predictable by the exchange rate. CRR argue that commodity prices can be thought of as the 'fundamentals' for the exchange rates of commodity exporters, implying that commodity prices should be predictable by commodity currencies. CRR test this 'commodity currency hypothesis' using currencies and country-specific commodity-price indices

¹Various predictors for commodity prices have been proposed in the literature, such as commodity forward prices (e.g., Fama and French, 1987; Gorton et al., 2013; Chinn and Coibon, 2014). Out-of-sample results remain however mixed, as remarked by Bernanke (2008), who emphasises that the unpredictability of commodity prices poses a major challenge to monetary policy.

for five commodity exporters (Australia, New Zealand, Canada, Chile and South Africa) and demonstrate that the currencies indeed hold predictive power for the commodity indices.

The commodity currency hypothesis does not necessarily imply that the predictive relation between exchange rates and commodity prices is linear. Moreover, additional variables may play a role. CRR (2010) therefore conclude their article by suggesting to study the robustness of their results to alternative nonlinear model specifications and omitted variables as a direction for future research. This is not a straightforward exercise because of the degrees of freedom involved. Nonlinear econometric models come in many forms, not to mention the sheer amount of other potential predictors. In predicting commodity prices, these nonlinearities may include, among others, regime switches (Mamatzakis and Remoundos, 2011; Beckmann and Czudaj, 2014; Chevallier et al., 2014), periods of booms and busts (Cashin, McDermott and Scott, 2002) and changes in the persistency (even local trends and explosive behavior) of commodity price levels (Kellard and Wohar, 2006; Gronwald, 2016). The relation between commodity prices and various other macroeconomic variables has also been studied extensively. Besides exchange rates, these variables include for example interest rates, industrial production (real activity), money and inflation (consumer prices) (e.g., Browne and Cronin, 2010; Akram, 2010), and oil prices (e.g., Wang et al., 2014; Ahmadi et al., 2016).² However, Pindyck and Rotemberg (1990), among others, argue that such fundamental variables do not fully explain the observed dynamics of commodity prices. Commodities are often treated as an investment class, rather than a production input. As with the prices of other financial assets, unobservable factors (e.g., investor psychology and heterogeneous expectations) play a role in driving commodity prices, leading to bubble-type patterns and excess volatility, thereby weakening the relation between commodity prices and its fundamentals (see, e.g., Arezki et al (2014) for a recent survey on the 'financialization' of commodity markets).

In addition to the choice of the correct nonlinear and/or multivariate model specification

²See also, e.g., Gospodinov and Ng (2013), Chen et al. (2014) and West and Wong (2014) for recent applications of factor models for commodity prices.

being ambiguous, commonly used nonlinear and multivariate regression models contain more parameters than the simple linear predictive regressions used by CRR. These additional parameters are costly to estimate in small samples, leading possibly to inferior forecasting ability of the nonlinear and multivariate model. In this study, we aim to tackle these issues by reconsidering the out-of-sample predictability documented by CRR in the context of so-called noncausal (vector) autoregressions, which are autoregressive models that parsimoniously allow for dependence on both future and past observations (see Brockwell and Davis, 1987; Breidt et al., 1991; Lanne and Saikkonen, 2011, 2013). Noncausal autoregressions can accommodate various nonlinearities and omitted variables, missed by conventional predictive models, without explicit specification, while containing the same number of parameters as simple linear causal autoregressions. The ambiguity regarding the correct nonlinear function form of commodity pricing models motivates the application of noncausal models in this context. For example, Deaton and Laroque (1992) propose a rational expectations competitive storage model with a central feature that the market as a whole cannot carry negative inventories, introducing non-linearity to the predicted commodity price series. However, they also acknowledge that their model does not yield a fully satisfactory explanation for nonlinearities and autocorrelation structure of actual commodity prices leaving out some unmodeled dynamics.

As we further discuss in Section 2, noncausal autoregressions have been recently applied successfully in modeling and forecasting various financial and macroeconomic variables. In particular in the presence of omitted variables or nonlinearities, noncausal models are found to fit the data better than their causal counterparts (see, e.g., Lanne and Saikkonen, 2011; Lanne et al. 2012b; Lof, 2013; Hencic and Gouriéroux, 2015). Gouriéroux and Zakoian (2016) show that noncausal models have an observationally equivalent nonlinear causal representation. We illustrate this relation between nonlinearity and noncausality with a small-scale simulation study in the Appendix. Based on these theoretical and empirical results, we believe that noncausal autoregressions are appropriate tools for investigating the

robustness of the commodity currency hypothesis to nonlinearities and omitted variables.

Using the same data as CRR (2014), we start by considering the predictability of commodity prices using conventional univariate and bivariate causal linear autoregressions. Our results confirm the main findings of CRR: Including exchange rates in the information set leads, for a number of countries, to more accurate out-of-sample forecasts of commodity price indices. After we expand the exercise to include noncausal models, we find that noncausal models in general do a better job at predicting commodity prices than their causal counterparts. Nevertheless, within the class of noncausal models we find less evidence that conditioning on exchange rates improves the out-of-sample predictability of commodity prices. Finally, we pool the forecasts across countries and concentrate on the qualitative differences between forecasts to increase the number of observations and gain statistical power. The results of this pooling exercise confirm that the increased forecasting accuracy from allowing for noncausality is statistically significant, while the incremental ability of exchange rates to predict commodity prices is insignificant.

We have also explored the predictability in the reverse direction, from commodity prices to exchange rates. Both causal and noncausal autoregressive models turn out to perform poorly at forecasting exchange rates. This should not come as a surprise as it is well known from the literature (including CRR) that it is hard to beat a random walk when it comes to forecasting exchange rates. In this paper, we therefore only report results on forecasting commodity prices.

Our results are consistent with those of Bork et al. (2014), who cast doubt on the commodity currency hypothesis by arguing that the predictive relations implied by the present value model (Eq. (1)) do not hold when the fundamentals f_t are themselves set by forward-looking financial markets. Market efficiency implies that tradeable assets should not be predictable by lagged information. As commodity prices and exchange rates simultaneously absorb all available and relevant information at the time of release, there can be no intertemporal predictability from one to another. Consistent with this view, Bork et

al. (2014) provide evidence that the relation between commodity prices and exchange rates is mainly contemporaneous, while the evidence for predictability is rather minor and not robust.

Although our paper also studies the robustness of the currency-commodity predictability, we look at the issue from a different angle than Bork et al. (2014). They attribute the discrepancies between their results and those of CRR mainly to data choices. The predictability is weakest when disaggregated individual commodity prices and end-of-period prices are used, instead of the period-averaged commodity indices used by CRR. Rather than evaluating the robustness by alternating the *data*, we focus on applying different predictive *models*. Even when using the exact same data as CRR, we show that the predictive power of exchange rates largely disappears with the alternative noncausal model.

The rest of the paper is organized as follows. In the next section (Section 2), we review noncausal autoregressions and discuss their merits for modeling commodity prices. In Section 3, we document out-of-sample forecasting performance of noncausal and various alternative linear and nonlinear models and forecast combinations. Finally, Section 4 concludes.

2 Methodology and data

2.1 Noncausal autoregression

In a noncausal autoregression, the variable of interest is allowed to depend both on past and future observations. Lanne and Saikkonen (2011) formulate a univariate noncausal autoregressive NCAR(r, s) process for the time-series y_t , depending on r lags and s leads:

$$\phi(L) \varphi(L^{-1}) y_t = \varepsilon_t, \quad (2)$$

where $\varphi(L^{-1}) = 1 - \varphi_1 L^{-1} - \dots - \varphi_s L^{-s}$, $\phi(L) = 1 - \phi_1 L - \dots - \phi_r L^r$, and L is the usual backshift operator (i.e., $L^k y_t = y_{t-k}$). The polynomials $\phi(z)$ and $\varphi(z)$ are assumed to have their zeros outside the unit circle. Furthermore, ε_t is an independently and identically distributed (*i.i.d.*) non-Gaussian error term.³ An intercept term is omitted from equation (2) for ease of notation but the possible nonzero mean of y_t is naturally taken into account throughout the analysis when computing fitted values and forecasts with noncausal models. When $s = 0$, Eq. (2) reduces to a conventional causal AR(r) model. On the other hand, when $r = 0$, the resulting NCAR($0, s$) model is purely noncausal with y_t depending only on its future values. When y_t is a vector rather than a scalar, Eq. (2) defines a noncausal vector autoregressive NCVAR(r, s) process (Lanne and Saikkonen, 2013).

Rewriting (2), it is easy to demonstrate that a noncausal time series depends on future error terms:

$$y_t = \phi_1 y_{t-1} + \dots + \phi_r y_{t-r} + v_t, \quad (3)$$

where

$$v_t = \varphi(L^{-1})^{-1} \varepsilon_t = \sum_{j=0}^{\infty} \beta_j \varepsilon_{t+j}, \quad (4)$$

in which the coefficients β_j can be recursively solved from the estimated noncausal part of the model, i.e. the parameters included in $\varphi(L^{-1})$ (see Lanne and Saikkonen, 2011, for details). The dependence of y_t on future error terms ε_{t+j} , implies that the errors are *nonfundamental*, i.e. they can not be interpreted as true economic shocks. Nonfundamentality arises when the information set of the agents in the economy is larger than the information set in the econometric model (see Hansen and Sargent, 1991; Alessi, 2011). Lanne and Saikkonen (2011) suggest that noncausal autoregressions take this missing information into account, by allowing for predictable errors. Lof (2013) carries out a simulation study to test this proposition. It appears indeed that when univariate autoregressions are applied to time

³As pointed out by Breidt et al. (1991), causal and noncausal autoregressive processes are indistinguishable when the error term is Gaussian. Following Lanne and Saikkonen (2011, 2013), we assume that the error terms in the noncausal models are t -distributed throughout this paper.

series variables that are actually generated as part of a multivariate system, noncausal autoregressions often fit the data better than causal autoregressions. Lof (2013) also finds that noncausal autoregressions tend to outperform linear causal models for variables that are generated by a nonlinear data generating process. Therefore, since CRR express their concerns about the role of omitted variables and nonlinearities on the predictive ability of exchange rates for commodity prices, it is worth considering how the predictive power holds up after using noncausal autoregressions, which is exactly what we do this paper.

Various recent studies point out the importance of addressing noncausality and nonfundamentality of financial variables. Kasa et al. (2014) consider a financial market in which agents receive private signals that are not observable to an econometrician observing market outcomes, and argue that this nonfundamentality can account for various empirical puzzles found in asset pricing. Lof (2013) shows empirically that S&P 500 returns and dividend yields are noncausal, and demonstrates by simulation that asset pricing models with heterogeneous speculators can generate noncausal representations of asset prices. Gouriéroux and Zakoian (2016) show that noncausal processes can capture the bubble-type patterns and other asymmetries often observed in financial markets, including commodity prices.

Figure 1: Impulse-response path of causal autoregression (left in figure (a): $y_t = \rho y_{t-1} + \varepsilon_t$) and noncausal autoregression (right in (b): $y_t = \rho y_{t+1} + \varepsilon_t$), with $\varepsilon_t = 1$ when $t = 0$ and $\varepsilon_t = 0$ otherwise. The autoregressive parameter ρ is set to $\rho = 0.5$.

To illustrate this point, Figure 1 displays the impulse-response path of a first-order causal and noncausal autoregressive process before and after a shock at $t = 0$. The causal model (left) shows the usual case of an unexpected jump that subsequently fades away gradually. For the noncausal model, the timing of events is reversed: The shock is anticipated, causing

the variable to inflate followed by a collapse at $t = 0$. This example clearly illustrates that the error term ε_t should not be interpreted as an 'economic shock', because the shock is anticipated.⁴ Moreover, the sign of ε_t at $t = 0$ is positive, while the plot displays a crash at $t = 0$, which should be interpreted as a negative shock in economic terms. This pattern of exponential growth followed by a collapse is defined by White and Granger (2011) as a bubble.

As commodity prices often display similar bubble-type patterns as in Figure 1 (see, e.g., Gutierrez, 2013; Etienne et al., 2014; Brooks, Prokopczuk and Wu, 2015), it could be expected that noncausal models may provide a good fit to commodity prices. Indeed, Karapanagiotidis (2014) and Gouriéroux and Jasiak (2016) find that the prices of commodity futures are well described by noncausal autoregressions. Also for the commodity indices in our dataset, we find strong evidence of noncausality.

2.2 Data and descriptive results

All results in this paper are computed using the dataset of CRR (2014), available on the website of Barbara Rossi. Throughout the analysis, we use the log-differences of the commodity prices and exchange rates (i.e. we are forecasting price changes). The dataset contains quarterly observations on country-specific commodity-price indices (cp_t) and exchange rates (s_t ; denominated in dollars) for five commodity exporters: Australia (AUS), New Zealand (NZ), Canada (CAN), Chile (CHI) and South Africa (SA). The country-specific commodity indices consist of various commodities, including metals, agricultural products and fossil fuels. Descriptive statistics of log-differences in the commodity price indices are presented in Table 1, showing clear non-Gaussian patterns with excess kurtosis and negative skewness as typically found in asset returns.

To quantify the ample evidence of noncausality (over the full sample), Table 2 shows the

⁴See Lanne and Nyberg (2015) for a more detailed discussion on impulse response analysis in the context of noncausal models.

Table 1: Descriptive statistics

Country	AUS	NZ	CAN	CHI	SA
Start date	1984Q1	1987Q1	1973Q1	1989Q3	1994Q1
Observations (T)	118	106	162	96	78
Mean	0.926	0.719	1.039	0.983	1.497
Std.dev.	5.491	5.141	7.985	15.171	8.056
Skewness	-0.132	-0.724	-2.240	-1.312	-1.691
Kurtosis	5.256	7.255	18.000	8.984	9.232
Jarque-Bera	25.355***	89.243***	1654.000***	170.764***	163.406***

Notes: This table presents the date of the first available observation, the number of observations for each country and descriptive statistics of the log-differenced prices of country-specific commodity baskets (multiplied by 100, i.e. percentages). For the Jarque-Bera normality test, *** denotes 1% statistical significance level. The last observation is in all cases 2013Q3.

values of the estimated log-likelihood functions for the causal AR(1) model and noncausal NCAR(0,1) model, applied to the log-differenced commodity indices of the five countries. The noncausal models yield higher full-sample log-likelihoods than the causal models for all countries, meaning that the noncausal model provide superior in-sample fit⁵.

Table 2: Estimated log-likelihood values for causal and noncausal models

Country	AUS	NZ	CAN	CHI	SA
AR(1)	-346.749	-310.489	-558.067	-387.993	-266.317
Jarque-Bera (residuals)	89.657***	59.487***	1145.791***	135.108***	160.454***
NCAR(0,1)	-338.938	-303.658	-529.714	-377.844	-258.604

Notes: This table reports the values of the full sample estimated log-likelihood functions of the AR(1) and NCAR(0,1) models fitted to the log-differenced commodity indices (Δcp_t) in different countries. The Jarque-Bera normality tests of the residuals from the causal AR(1) model are also presented.

Concerning the adequacy of the causal AR(1) model in terms of residual diagnostics, Table 2 clearly indicates considerable deviation from normality: The p -values of the Jarque-Bera test are less than 0.001 for each case. A more leptokurtic distribution, such as the t -distribution is therefore suitable for this data. This non-normality, also clearly present in Table 1, is the necessary first step for the use of noncausal models based on the t -distribution, which appears an adequate distributional assumption for our quarterly data.

⁵Lanne and Saikkonen (2011, 2013) propose to select the optimal model by comparing the log-likelihood of the causal AR(p) and noncausal NCAR(r, s) models of the same order (i.e. $r + s = p$). Throughout this paper, we follow CRR in reporting only results on first-order models (i.e. the AR(1) and NCAR(0,1) models). As a robustness check, we also consider order-selection by Akaike (AIC) and Bayesian (BIC) information criteria. The BIC typically selects a first-order model, such that the results are near identical to those reported in this paper. The AIC often selects higher orders, but these models generally underperform out of sample compared to the first-order benchmarks. Results are available upon request.

As a representative example, in Figure 2 we depict the residual quantile-quantile (Q-Q) plots for New Zealand based on the estimated causal AR(1) and the noncausal (NCAR(0,1)) models, whose estimated values of the log-likelihood functions are reported in Table 2. The right panel of Figure 2 lends support to the adequacy of the t -distribution of the errors. In particular, the t -distribution seems to capture the tails reasonably well. Moreover, the estimate of the degree-of-freedom parameter in the NCAR(0,1) model is rather small 3.96, suggesting the inadequacy of the Gaussian error distribution.

Figure 2. Quantile-quantile (Q-Q) plots of the residuals of the Gaussian AR(1) (left panel) and NCAR(0,1) (right panel, based on the t -distribution) models for New Zealand. Data quantiles are on the y-axis while theoretical quantiles from the normal distribution (left) and the t -distribution (right) are depicted on the x-axis.

2.3 Forecasting with noncausal models

In Section 3.1, we demonstrate that, in addition to in-sample predictive performance, the noncausal models also perform better out of sample. Before that, we review the forecasting routine for noncausal autoregressions. As the prediction problem in the univariate noncausal AR model is generally nonlinear and no analytical solution is available, we rely on the simulation-based method by Lanne et al. (2012a). Using the representation (Eq. (3)), the mean-square sense optimal one-step-ahead forecast is the conditional expectation, given the

information set at time T :

$$E_T(y_{T+1}) = \phi_1 y_T + \dots + \phi_r y_{T+1-r} + E_T(v_{T+1}). \quad (5)$$

The information set includes the data y_t up to T ($t = 1, \dots, T$), which is also the available information in parameter estimation. The forecast (Eq. (5)) requires a forecast of the noncausal (forward-looking) component $v_{T+1} = \sum_{j=0}^{\infty} \beta_j \varepsilon_{t+1+j}$ (see Eq. (4)). Following Lanne et al. (2012a), the forecast of v_{T+1} is based on the approximation

$$v_{T+1} \approx \sum_{j=0}^{M-1} \beta_j \varepsilon_{T+1+j}, \quad (6)$$

where the integer M is assumed to be large enough to make the approximation error negligible. Notice that the truncated sum (Eq. (6)) depends on the future error terms ε_{T+1} , ε_{T+2} up to ε_{T+M} . To a close approximation, we have then

$$E_T(y_{T+1}) \approx \phi_1 y_T + \dots + \phi_r y_{T+1-r} + E_T\left(\sum_{j=0}^{M-1} \beta_j \varepsilon_{T+1+j}\right). \quad (7)$$

Lanne et al. (2012a) derive the distribution of $(\varepsilon_{T+1}, \dots, \varepsilon_{T+M})$, conditional on the observed data $(y_1, \dots, y_T)$. From this conditional distribution we simulate R mutually independent realizations of $(\varepsilon_{T+1}, \dots, \varepsilon_{T+M})$, after which their average is plugged into Eq. (7) to obtain the one-period-ahead forecast of y_{T+1} .⁶ The fact that only data observed at time T is used as the conditional information in the forecasting routine above emphasizes the important point that despite the forward-looking nature of the noncausal model, this forecasting procedure does not suffer a look-ahead bias. Hence, the forecasts obtained with the noncausal model can be compared to the causal model.

The conditional expectation $E_t(y_{t+1})$ derived from the noncausal model is in general a

⁶As Lanne et al. (2012) and Nyberg and Saikkonen (2014), we set $R=10,000$ for univariate models, $R=200,000$ for multivariate models, and $M=50$ throughout this paper.

nonlinear function of lagged values of y_t , although a closed form expression of the conditional distribution and expectation is generally not available (see also Gouriou and Zakoian, 2016). The nonlinearity implied by the noncausal model becomes clear from Figure 3, which shows scatterplots of y_t against the conditional expectation $E_t(y_{t+1})$, derived from the full-sample estimates of both the causal and noncausal models in Table 2 and for different histories t .⁷ The conditional expectations derived from the causal AR(1) model clearly line up linearly, because $E_t(y_{t+1}) = \phi_1 y_t$. The conditional expectations derived from the noncausal NCAR(0,1) model are computed by simulation, following Eq. (7). Because $r = 0$, the backward looking AR parameters $\phi_1, \dots, \phi_r$ drop out of the NCAR(0,1) model such that the conditional expectation reduces to

$$E_t(y_{t+1}) = E_t\left(\sum_{j=0}^{M-1} \beta_j \varepsilon_{t+1+j}\right). \quad (8)$$

Unlike with the causal model, the noncausal conditional expectations form an S-shaped curve, which mitigates the expected persistence of outliers. Figure 3 hence illustrates one of the main advantages of noncausal autoregressions: We are able to fit a nonlinear curve to the data, while the model does not contain any more parameters to estimate than a standard causal autoregression. In the Appendix, we illustrate this point further with a simulation experiment, by showing that time series generated by noncausal models clearly display nonlinear dependence. Moreover, in Section 3 we show that this greater flexibility of noncausal models leads to more accurate out-of-sample forecasts.

To compute forecasts with the noncausal VAR, we use the methodology proposed by Nyberg and Saikkonen (2014), who extend the ideas explained above to a multivariate setting (see details in their paper). Lanne et al. (2012a) and Nyberg and Saikkonen (2014) apply these methods to inflation forecasting and find that the noncausal models achieve superior forecasting performance compared to causal models. Moreover, Lanne et al. (2012b)

⁷In the forecasting exercise (Section 3), we use the usual expanding and rolling window approaches, instead of the full sample, to compute (out-of-sample) forecasts.

Figure 3: Scatterplots of y_t against $E_t(y_{t+1})$, computed using the full-sample parameter estimates (see Table 1) of the causal AR(1) model (grey dots) and noncausal NCAR(0,1) models (black dots, see Eq. (8)).

compare the forecasting performance of (univariate) noncausal and causal autoregressive models for a large monthly and quarterly dataset of the U.S. macroeconomic and financial time series. The noncausal models, accommodating omitted variables and moderate conditional heteroskedasticity, consistently outperform causal models especially for quarterly time series. The quarterly data of interest in this study is the frequency for which these advantageous features of noncausal models are expected to be the greatest. For higher frequency data, in line with extensive amount of evidence in empirical finance research, non-linear methods incorporating, e.g., GARCH components are generally needed to adequately address the strong conditional heteroskedasticity present in asset returns.

3 Forecasting results

We consider the out-of-sample forecasting performance of causal and noncausal autoregressive models in several stages. In Section 3.1, we begin by replicating the out-of-sample forecasting results of CRR (2014) using standard (causal) predictive regressions. After that, we reconsider forecasting results based on noncausal autoregressions. In addition to these

analyses, in Section 3.2 we present various robustness checks to our main forecasting results.

3.1 Out-of-sample forecasting performance

Following CRR (2010, 2014), for each country we first consider a (causal) AR(1) model for the log-differenced commodity prices:

$$\Delta cp_t = \gamma_0 + \gamma_1 \Delta cp_{t-1} + \eta_t, \quad (9)$$

where the (possible) predictive power is coming solely from the past lagged changes in the commodity prices. Next, we examine the forecasting performance of exchange rates by considering a bivariate VAR(1) model for log-differenced commodity prices and exchange rates. The first equation of the VAR is equivalent to the predictive model considered by CRR, allowing for dependence on lagged exchange rates:

$$\Delta cp_t = \gamma_0 + \gamma_1 \Delta cp_{t-1} + \gamma_2 \Delta s_{t-1} + \eta_t. \quad (10)$$

If the commodity currency hypothesis is valid, the additional (out-of-sample) predictive power coming from the inclusion of exchange rates (Δs_{t-1}) is expected to lead superior forecasts over the univariate AR(1) model (Eq. (9)).

Results on the predictive power of exchange rates in the conventional context of causal predictive regressions (Eq. (9) and (10)) are presented in Table 3. In accordance with CRR (2010, 2014), the first row contains (in-sample) Granger causality test statistics, computed from the VAR model. We find evidence that exchange rates Granger cause commodity prices for three countries (Australia, New Zealand, and Chile). When we look at the out-of-sample forecasting performance, South Africa can be added to this list as well: The second row of Table 3 shows the relative out-of-sample Mean Squared Forecast Errors (MSFEs) of a VAR(1) (Eq. (10)) relative to the MSFE of an AR(1) (Eq. (9)). Similar to CRR, we compute

one-quarter-ahead out-of-sample forecasts for the second half of the sample separately for each country using an expanding window approach, in which the first estimation sample is the first half of the full sample ($0.5T$).⁸

Table 3: Causal predictive regressions

Country	AUS	NZ	CAN	CHI	SA
Granger causality	3.914**	3.727**	0.004	3.608*	1.235
Relative MSFE	0.943	0.945	1.012	0.973	0.998
Relative MAFE	0.987	0.977	1.004	0.981	0.989

Notes: In-sample Granger causality test statistics ($H_0 : \Delta s$ does not Granger cause Δcp) obtained from a VAR(1) fitted to Δcp_t and Δs_t . The relative out-of-sample mean-squared forecast errors (MSFE) and mean-absolute forecast errors (MAFE) evaluate the forecasts of Δcp using a bivariate VAR(1) (including Δcp_t and Δs_t) relative to a univariate AR(1) (including only Δcp_t). Asterisks indicate 10% (*), 5% (**) and 1% (***) significance levels. Granger causality tests are based on heteroskedasticity-autocorrelation consistent Wald test statistics.

Our results are consistent with the out-of-sample results by CRR (2014, Table IV): For all countries but Canada we find that exchange rates have predictive power, as the VAR(1) model produces more accurate forecasts (i.e. smaller MSFEs) than the univariate AR(1) model. However, the results of Diebold-Mariano (1995) and West (1996) test reveal that the differences in forecast errors between the AR and VAR models are not statistically significant at conventional significance levels⁹. Similar evidence is obtained with the Mean Absolute Forecast Errors (MAFE) (see the last row of the table) which is a more robust measure for the possible effects of large outliers than the MSFE. This provides additional robustness to our findings with respect of alternative loss functions than the usual MSFE criterion. We suspect that the lack of statistical significance is largely due to the limited number of observations available per country (reported in Table 1). Later in this section, we therefore pool the predictions of each model across countries and consider qualitative differences (rather than the MSFE and MAFE loss functions) between the models. Before that, we turn to the forecasts of the noncausal models.

⁸As a robustness check, we also apply a rolling window approach, with very similar results (available upon request).

⁹Like CRR, we do find significantly smaller MSFEs when we compare the VAR to a Random Walk (i.e. the VAR model, containing the effect of exchange rates, outperforms the Random Walk). For ease of notation, we only report the comparison between the AR and VAR models here.

Table 4 reports the forecasting accuracy of noncausal models relative to the causal AR(1) (see (Eq. 9)). First, we obtain forecasts for commodity prices using a first-order univariate noncausal autoregression NCAR(0,1). The relative MSFE for the NCAR(0,1) is reported for each country in the first row of Table 4. We proceed by testing the predictive power of exchange rates within the noncausal framework by forecasting commodity price changes using a bivariate noncausal vector autoregression NCVAR(0,1) that also includes exchange rates (both variables are measured in log-differences, as for the causal models). Overall, the application of noncausal models seems to deliver some improvement to the forecasting performance. Although not statistically significant, the NCAR(0,1) yields smaller MSFEs than the causal AR(1) for most countries. For the vector autoregressions, the results are mixed: The NCVAR(0,1) does not perform strictly better or worse than the univariate NCAR(0,1) or the causal AR(1) and VAR(1) across countries.

Table 4: Noncausal forecasting results

Country	AUS	NZ	CAN	CHI	SA
NCAR(0,1)	0.993	1.008	0.975	0.898	0.999
NCVAR(0,1)	0.931*	0.993	0.963	1.006	1.023
Forecast combinations					
Univariate: AR(1) & NCAR(0,1)	0.989	0.994	0.978	0.937	0.998
Multivariate: VAR(1) & NCVAR(0,1)	0.930*	0.953	0.979	0.982	1.009
Causal: AR(1) & VAR(1)	0.961*	0.967	1.006	0.980	0.998
Noncausal: NCAR(0,1) & NCVAR(0,1)	0.949*	0.997	0.968	0.938	1.008

Notes: Relative out-of-sample MSFEs obtained from forecasts of Δcp using the model in the first column relative to an AR(1) model. The forecasts combinations (last 4 rows) apply equal weights to each model. Asterisks indicate 10% (*), 5% (**) and 1% (***) significance levels, based on a Diebold-Mariano (1995) and West (1996) test.

The final rows of Table 4 show various forecast combinations. First, we combine the forecasts of causal and noncausal univariate models: AR(1) and NCAR(0,1), which we compare to the forecast combination of the multivariate VAR(1) and NCVAR(0,1) models with an explicit dependence on exchange rates. We apply equal weights to the models, which are typically considered optimal.¹⁰ Next, we combine the forecasts from the causal AR(1) and VAR(1) models, which we compare to the combination of noncausal models. Although

¹⁰See, e.g., Baumeister et al. (2014), for an application that involves pooling various forecasting models for oil prices; and Timmermann (2006).

these forecast combinations yield overall improvements with respect to the individual models, in some cases even statistically significant, a clear pattern is still lacking: No forecast combination dominates another combination across countries.

In Table 5 such a pattern does emerge. We pool the forecasts over all countries for each model, and run horse races between different models. Table 5 reports the fraction of observations for which a given model (A) yields a smaller squared forecast error than another model (B). That is, we consider whether model (A) qualitatively outperforms model (B). We test whether the fraction deviates significantly from 50% using a Diebold-Mariano (1995) sign test. The results are reported for both expanding and rolling window approaches.

Table 5: Pooled out-of-sample predictability

Model A	Model B	Fraction	Fraction
		Exp. window	Rolling window
VAR(1)	AR(1)	0.468	0.521
NCAR(0,1)	AR(1)	0.525	0.539*
NCAR(0,1)	VAR(1)	0.479	0.500
NCVAR(0,1)	AR(1)	0.536	0.514
NCVAR(0,1)	VAR(1)	0.507	0.532
NCVAR(0,1)	NCAR(0,1)	0.504	0.492
Forecast combinations			
Multivariate	Univariate	0.514	0.504
Noncausal	Causal	0.546*	0.546*

Notes: Fraction denotes the percentage of observations (pooled over all countries) for which model (A) yields more accurate out-of-sample forecast of Δcp (i.e. smaller squared forecast error) than model (B). Asterisks indicate whether fractions exceed 50% (H_0 : Fraction=0.5) at the 10% (*), 5% (**) and 1% (***) levels, based on a Diebold-Mariano (1995) sign test. Following CRR (2010, 2014), the size of the rolling window equals 50% of the full sample.

From Table 5, we can see that noncausal NCAR and NCVAR models yield superior forecasts (i.e. provide more accurate forecasts for more than 50% of the observations) compared to their causal AR and VAR counterparts. The multivariate VAR and NCVAR models that include exchange rates, however, do not yield significantly more accurate forecasts than the univariate AR and NCAR models. Instead of conditioning on exchange rates, it seems that allowing for noncausality is a more fruitful approach for predicting commodity prices. The forecast combinations, especially, show this pattern: The noncausal model combination yields more accurate forecasts than the causal combination and this difference in forecasting

performance is also statistically significant. In contrast, the difference between the multivariate and univariate combinations is not statistically significant.

3.2 Robustness checks

In this section, we verify the robustness of the out-of-sample forecasting results in Section 3.1, by considering absolute forecast errors rather than squared forecast errors, by evaluating time-variation in predictability, and by comparing our results to commonly used nonlinear autoregressive models (i.e. Smooth Transition Auto Regressive (STAR) Models).

The forecast evaluation in Table 4 is based on mean squared forecast errors (MSFE). In Table 3, we already compare the forecast performance of causal predictive regressions using both the MSFE criterion and mean absolute forecast errors (MAFE). In Table 6, we report the forecast comparisons using the MAFE criterion for the same models and forecast combinations as in Table 4. As far as the relative performance of the noncausal models and the predictive power of exchange rates is concerned, the results are similar to those in Table 4. While the relative MAFE statistics are slightly smaller than the MSFEs in Table 4, the general findings appear robust to the choice of the forecast loss criterion.

Table 6: Noncausal forecasting results (relative MAFEs)					
Country	AUS	NZ	CAN	CHI	SA
NCAR(0,1)	0.990	1.007	0.984	0.963	1.005
NCVAR(0,1)	0.969	1.001	0.971	0.994	1.016
Forecast combinations					
Univariate: AR(1) & NCAR(0,1)	0.987	0.998	0.991	0.981	1.003
Multivariate: VAR(1) & NCVAR(0,1)	0.976	0.971	0.988	0.986	1.001
Causal: AR(1) & VAR(1)	0.984	0.983*	1.002	0.990	0.992
Noncausal: NCAR(0,1) & NCVAR(0,1)	0.970	0.998	0.977	0.975	1.008

Notes: Relative out-of-sample MAFEs obtained from forecasts of Δcp using the model in the first column relative to an AR(1) model. See also the notes to Table 4.

It is worth noting that when considering the qualitative differences between different forecasts in the same spirit as presented in Table 5, it does not matter whether mean-squared or mean-absolute forecast errors are used as the measure of forecast accuracy: The ranking between the models and forecasts is equivalent for these two criteria. Therefore,

the results measured by the fractions are the same for both MSFE and MAFE loss criteria.

Figure 4: Black bars show the fraction of quarterly observations in each year (pooled over all countries) for which the forecast combination of noncausal models yields more accurate forecasts than the forecast combination of causal models. Grey bars show the fraction of observations for which the combination of multivariate models yields more accurate forecasts than the combination of univariate models.

It is meaningful to examine possible time-variation in the predictive performance of the different competing models, as it is well-documented in the (financial) econometric forecasting literature that predictability often varies over time. Table 5 in the previous section shows that the forecast combination of noncausal models yields over the full sample more accurate out-of-sample forecasts than the combination of causal models. The black bars in Figure 4 present this forecast comparison separately for each year, starting from 2004 when we obtain forecasts for all five countries in the sample. In each year, there are twenty observations (five countries, four quarters) for which we compute forecasts. The black bar shows the fraction of observations for which the noncausal forecast combination (NCAR(0,1) and NCVAR(0,1)) yields more accurate forecasts than the causal forecast combination (AR(1) and VAR(1)) (cf., e.g., Table 5). This fraction is fairly stable and in the 10 years under consideration, the combined forecasts from noncausal models strictly outperform the causal models during 8 years, while the causal models outperform the noncausal models only once, in 2008. The grey bars show the comparison of multivariate (VAR(1) and NCVAR(0,1)) and univariate (AR(1) and NCAR(0,1)) forecast combinations in each year. This pattern is

clearly more erratic, compared to the black bars. The improved forecasting performance of noncausal models over causal models is therefore rather resilient, while the relative performance of multivariate models over univariate models heavily fluctuates across years. The results in Figure 4 therefore contradict the commodity currency hypothesis, since conditioning on exchange rates does not result in consistently more accurate forecasts of commodity prices.

Finally, for comparison, we compute forecasts of changes in commodity prices using a nonlinear Logistic Smooth Transition Auto Regressive (LSTAR; see, e.g., Teräsvirta, 1994), where an explicit nonlinear form of the model is employed instead of the noncausal approach considered above. LSTAR models have been applied in the literature to model financial time series, including exchange rates (e.g., Taylor et al., 1999) as well as commodity prices. For example, Beckmann and Czudaj (2014) apply an LSTAR model to examine the relationship between prices for first and second nearby futures contracts of seven agricultural commodities. Their specification allows for so called contango (futures price is above the spot price) and backwardation regimes, which are neglected in a linear framework.

We consider the following LSTAR specification to changes in commodity prices:

$$\begin{aligned}\Delta cp_t &= \theta_0 + \theta'_1 X_{t-1} G(z_{t-1}) + \theta'_2 X_{t-1} (1 - G(z_{t-1})) + \varepsilon_t \\ G(z_{t-1}) &= (1 + \exp(-\gamma z_{t-1}))^{-1},\end{aligned}\tag{11}$$

where the transition variable z_{t-1} is either the lagged change in commodity prices ($z_{t-1} = \Delta cp_{t-1}$) or the exchange rate ($z_{t-1} = \Delta s_{t-1}$) and X_{t-1} is either equal to lagged commodity price changes ($X_{t-1} = \Delta cp_{t-1}$), or it includes both commodity prices and exchange rates ($X_{t-1} = (\Delta cp_{t-1}, \Delta s_{t-1})$). That is, we consider four different specifications of the LSTAR model, where we allow for univariate nonlinear dynamics of Δcp_{t-1} , as well as for both linear and nonlinear dependence on lagged exchange rates to examine the commodity currency hypothesis.¹¹

¹¹In addition, we also consider a STAR model with an exponential transition function $G(z_{t-1})$, but find

Table 7: LSTAR forecasting results

Panel A: Relative MSFE						
z_{t-1}	X_{t-1}	AUS	NZ	CAN	CHI	SA
Δcp_{t-1}	Δcp_{t-1}	1.028	1.059	1.086	1.026	1.008
Δcp_{t-1}	$(\Delta cp_{t-1}, \Delta s_{t-1})$	1.142	1.068	1.102	1.059	1.027
Δs_{t-1}	Δcp_{t-1}	1.061	1.025	1.398	0.937	1.055
Δs_{t-1}	$(\Delta cp_{t-1}, \Delta s_{t-1})$	1.069	1.041	1.895	1.598	1.055

Panel B: Relative MAFE						
z_{t-1}	X_{t-1}	AUS	NZ	CAN	CHI	SA
Δcp_{t-1}	Δcp_{t-1}	0.975	1.059	1.038	0.949	1.013
Δcp_{t-1}	$(\Delta cp_{t-1}, \Delta s_{t-1})$	1.029	1.058	1.045	0.968	1.051
Δs_{t-1}	Δcp_{t-1}	1.048	1.059	1.113	0.924	1.074
Δs_{t-1}	$(\Delta cp_{t-1}, \Delta s_{t-1})$	1.010	1.077	1.229	1.125	1.054

Notes: Out-of-sample MSFEs (panel A) and MAFEs (panel B) of the LSTAR model (Eq. (11)) with different specifications of X_{t-1} and z_{t-1} , relative to an AR(1) model. Asterisks indicate 10% (*), 5% (**) and 1% (***) significance levels, based on a Diebold-Mariano (1995) and West (1996) test.

We use the LSTAR model (Eq. (11)) for computing commodity price forecasts for each country in our sample with the same setup as in Section 3.1: That is, an expanding window approach, obtaining one-period ahead out-of-sample forecasts for the second half of the sample. Table 7 presents the results from this forecasting exercise. Panel A shows for each country the MSFEs of the various LSTAR specifications relative to the MSFEs of an AR(1) model, while panel B shows the relative MAFEs. Overall, the results suggest that the nonlinear LSTAR model yields less accurate forecasts than the simple causal AR(1) model, despite previous evidence of nonlinearity in commodity prices and exchange rates, and that exchange rates do not contain predictive power. None of the considered LSTAR specification significantly outperforms a simple linear AR(1) model. Moreover, these results do not lend support to the commodity currency hypothesis, since the models with explicit linear or nonlinear dependence on exchange rates (second to fourth row in Table 7) do not perform better than the univariate LSTAR model (first row). The poor forecasting performance of the LSTAR models is plausibly due to the fact that compared to the simple AR(1) model, the LSTAR model has two or four additional parameters to estimate, which are costly to

that this ESTAR model fits the data poorly, leading to imprecise parameter estimates and far inferior forecasts.

estimate in small samples.¹² Noncausal autoregressions, on the other hand, have the same number of parameters as causal autoregressions, which is appealing in particular in small samples, and turn out to lead more accurate forecasts than the STAR model.

4 Conclusions

Consistent with CRR (2010, 2014) and the commodity currency hypothesis, we find that exchange rates of several commodity exporters are useful for predicting commodity prices when applying linear causal predictive models. After expanding the class of models to include noncausal autoregressions, which may provide a parsimonious approximation to various nonlinearities and omitted predictors, we obtain slightly better forecasts overall, while the predictive power of exchange rates largely disappears.

Finding noncausal dynamics in commodity prices is consistent with recent research and implies that there may exist relevant predictors beyond the lagged values of commodity prices alone. In a noncausal autoregression, predictable future error terms act as a proxy for missing information, but the exact information set is unknown and could even include unobservables. While these unknown predictors may well include exchange rates, explicitly conditioning on exchange rates by adding them as predictors to a noncausal VAR, does not result in any measurable improvement in terms of forecasting accuracy.

CRR (2010) suggest to examine the robustness of their results by addressing the role of nonlinearities and alternative macroeconomic and financial predictors in future research. Our results confirm that more complex models are indeed fruitful: Commodity prices are predictable beyond a linear AR(1) model. Nevertheless, the results also suggest that this predictability of commodity prices does not come from exchange rates. Our results can therefore be interpreted as a rejection of the commodity currency hypothesis: Condition-

¹²The LSTAR model is often specified with an additional constant (scale) parameter in the transition function, which we normalized to zero. Without this normalization, the LSTAR forecasts are less accurate than the results presented in Table 7.

ing on exchange rates does not significantly improve the out-of-sample forecastability of commodity prices.

Our two main findings, noncausality of commodity prices and the limited predictive power of exchange rates, relate to recent research on increased speculative behavior on commodity markets and the 'financialization' of commodity markets in general (see Arezki et al., 2014, for an overview). Lof (2014) shows that the prevalence of speculative trading strategies can generate price dynamics that are well captured by noncausal autoregressions. Moreover, Hencic and Gouriéroux (2015) find noncausal dynamics in Bitcoin rates, which they attribute specifically to speculative online trading.

The process of 'financialization' that commodity markets experienced in recent decades means that commodities are being regarded as a financial asset class similar to stocks, bonds, exchange rates, or other securities (Arezki et al., 2014). Commodity prices are therefore expected to be highly liquid and instantly adjust to newly released information, similar to the prices of other financial assets traded on competitive financial markets. Easily predictable patterns in commodity prices are therefore subject to arbitrage trading. This leaves little predictive content to any lagged observable information, including exchange rates. Consistent with this implication, we find that predictability of commodity price changes is small in general - presumably too small to exploit for trading gains, and is not originating from observable exchange rates.

Appendix: Noncausality and nonlinearity

The noncausal autoregressive model (Eq. (2)) has an observationally equivalent causal (generally nonlinear) representation. Except for Gaussian noncausal models, the causal representation of a noncausal autoregression is nonlinear, and finding a closed-form solution for the causal representation of a noncausal autoregression is not always possible. Gouriéroux and Zakoian (2016) derive the nonlinear causal representation of a first-order noncausal autoregression with Cauchy distributed errors. Brockwell and Davis (1987, pp. 124–125)) and Lanne and Saikkonen (2013) show that Gaussian noncausal and causal (vector) autoregressions are observationally equivalent. Lof (2014) derives the explicit linear causal representation of the Gaussian noncausal VAR(1) model.

In this Appendix, we illustrate the nonlinearity of non-Gaussian noncausal autoregressions by a simulation experiment. We simulate realizations from the following first-order causal and noncausal autoregressive processes:

$$\begin{aligned} AR(1, 0) \quad y_t &= \phi y_{t-1} + \varepsilon_t \\ AR(0, 1) \quad y_t &= \varphi y_{t+1} + \varepsilon_t, \end{aligned} \tag{12}$$

where the error term ε_t are Gaussian (N ; standard normal), t -distributed (t_d ; with d degrees of freedom), or centered χ^2 -distributed (χ_d^2 ; with d degrees of freedom). We then use the simulated data to fit the causal model

$$y_t = \alpha_1 y_{t-1} + \alpha_2 y_{t-1}^2 + \alpha_3 y_{t-1}^3 + \epsilon_t, \tag{13}$$

which is the third-order Taylor approximation of a nonlinear first-order autoregressive model of the form $y_t = f(y_{t-1}; \beta) + \xi_t$, where $f(\cdot; \beta)$ is of at least third-order differentiability (e.g., a logistic STAR model). We test the hypothesis $H_0 : \alpha_2 = \alpha_3 = 0$ (implying there is no nonlinear dependence between y_t and y_{t-1}).

Table A1 reports the rejection frequencies after $R = 10,000$ simulation replications, with

Table A1: Simulation results										
Panel A: Small sample ($T=100$)										
	Causal					Noncausal				
level	N	t_3	t_8	χ_3^2	χ_8^2	N	t_3	t_8	χ_3^2	χ_8^2
10 %	0.09	0.07	0.08	0.08	0.08	0.09	0.62	0.22	0.46	0.26
5 %	0.04	0.04	0.04	0.04	0.04	0.04	0.55	0.15	0.36	0.18
1 %	0.01	0.01	0.01	0.01	0.01	0.01	0.40	0.06	0.21	0.08

Panel B: Large sample ($T=1000$)										
	Causal					Noncausal				
level	N	t_3	t_8	χ_3^2	χ_8^2	N	t_3	t_8	χ_3^2	χ_8^2
10 %	0.10	0.07	0.09	0.08	0.09	0.10	0.99	0.51	0.97	0.77
5 %	0.05	0.04	0.05	0.04	0.04	0.05	0.99	0.41	0.95	0.68
1 %	0.01	0.01	0.01	0.01	0.01	0.01	0.97	0.24	0.87	0.46

Notes: Rejection frequencies for the hypothesis $H_0 : \alpha_2 = \alpha_3 = 0$ (Eq. (13)) on time series of $T = 100$ or $T = 1000$ observations generated by a causal and noncausal first-order autoregression (Eq. (12)) with Standard Normal, t_d , or centered χ_d^2 -distributed errors ($d = 3$ or $d = 8$ refers to the degrees of freedom).

$\phi = \varphi = 0.8$, and $d = 3$ or $d = 8$ degrees of freedom. We initially set the simulated sample size to $T = 100$, of similar magnitude as the samples consider in this paper. We also examine large-sample behavior by setting $T = 1000$. The results clearly confirm that a noncausal autoregression with non-Gaussian errors implies nonlinear dependence between y_t and y_{t-1} : The rejection frequencies are close to the nominal significance levels only when evaluating the data generated by causal processed or by noncausal processes with Gaussian errors. When the test is applied to noncausal non-Gaussian data, the rejection frequencies are much higher. The rejection frequencies approach 100% when the simulated sample is large ($T = 1000$) and the underlying distribution is far from Gaussian ($d = 3$), but even for smaller samples and higher degrees of freedom, the rejection frequencies clearly exceed the nominal levels.

Figure A1 shows a single realization of $T = 1000$ simulated observations for both causal and noncausal time series with Gaussian, t_3 and χ_3^2 distributed errors. These scatterplots of y_{t-1} against y_t also show strong nonlinear patterns in the time-series generated by non-Gaussian noncausal autoregressions. For the causal time series, the fitted regression is essentially a straight line, even when the data is leptokurtic or skewed. For the noncausal time series, the regression line is straight only for the Normally distributed data. With the non-Gaussian noncausal time series, the data reveal a nonlinear curve.

Figure A1: Scatterplots of $T = 1000$ observations of y_{t-1} against y_t , simulated from a first-order causal (top row) and noncausal (bottom row) autoregression (Eq. (12)), with Normally distributed (left column), t -distributed (middle column), or centered χ^2 -distributed (right column) errors. The solid line shows the fitted values of a third order polynomial (Eq. (13)).

References

Ahmadi, M., Behmiri, N.B., Manera, M., 2016. How is volatility in commodity markets linked to oil price shocks? *Energy Economics* 59, 11–23.

Akram, F.Q., 2010. Commodity prices, interest rates and the dollar. *Energy Economics* 31, 838–851.

Alessi, L., Barigozzi, M., Capasso, M., 2011. Nonfundamentalness in structural econometric models: A review. *International Statistical Review* 79, 16–47.

Arezki, R., Loungani, P., van der Ploeg, R., Venables, A.J., 2014. Understanding international commodity price fluctuations. *Journal of International Money and Finance* 42, 1–8.

Baumeister C., Kilian, L., Lee, T.K., 2014. Are there gains from pooling real-time oil price forecasts? *Energy Economics* 46, Supplement 1, S33–S43.

Beckmann, J., Czudaj, R., 2014. Non-linearities in the relationship of agricultural futures prices. *European Review of Agricultural Economics* 41, 1–23.

Bernanke, B., 2008. Outstanding issues in the analysis of inflation. Federal Reserve Bank

of Boston 53rd Annual Economic Conference, <http://goo.gl/Uhh01Q> (accessed 23.12.2016)

Bork, L., Kaltwasser, P., Sercu, P., 2014. Do exchange rates really help forecasting commodity prices? SSRN working paper. (accessed at SSRN: <https://ssrn.com/abstract=2473624>).

Breidt, J., Davis, R.A., Lii, K.S., Rosenblatt, M., 1991. Maximum likelihood estimation for noncausal autoregressive processes. *Journal of Multivariate Analysis* 36, 175–198.

Brockwell, P.J., Davis, R.A., 1987. *Time Series: Theory and Methods*. Springer-Verlag.

Brooks, C., M. Prokopczuk, Wu, Y., 2015. Booms and busts in commodity markets: Bubbles or fundamentals? *Journal of Futures Markets* 35, 916–938.

Browne, F., Cronin, D., 2010. Commodity prices, money and inflation. *Journal of Economics and Business* 62, 331–345.

Campbell, J.Y., Shiller, R.J., 1987. Cointegration and tests of present value models. *Journal of Political Economy* 95, 1062–1088.

Cashin, P., McDermott, C.J., Scott, A., 2002. Booms and slumps in world commodity prices. *Journal of Development Economics* 69, 277–296.

Chen, S.L., Jackson, J.D., Kim, H., Resiandini, P., 2014. What drives commodity prices? *American Journal of Agricultural Economics* 96, 1455–1468.

Chen, Y., Rogoff, K., Rossi, B., 2010. Can exchange rates forecast commodity prices? *Quarterly Journal of Economics* 125, 1145–1194.

Chen, Y., Rogoff, K., Rossi, B., 2014. Can exchange rates forecast commodity prices? An update. Working paper available at: <http://www.econ.upf.edu/~brossi/>.

Chevallier, J., Gatumel, M., Ielpo, F., 2014. Commodity markets through the business cycle. *Quantitative Finance* 14, 1597–1618.

Chinn, M.D., Coibion, O., 2014. The predictive content of commodity futures. *Journal of Futures Markets* 34, 607–636.

Deaton, A., Laroque, G., 1992. On the behaviour of commodity prices. *Review of Economic Studies* 59, 1–23.

Diebold, F.X., R. Mariano, R., 1995. Comparing predictive accuracy. *Journal of Busi-*

ness and Economic Statistics 13, 134–144.

Engel, C., West, K.D., 2005. Exchange rates and fundamentals. *Journal of Political Economy* 113, 485–517.

Etienne, X.L., Irwin, S.H., Garcia, P., 2014. Bubbles in food commodity markets: Four decades of evidence. *Journal of International Money and Finance* 42, 129–155.

Fama, E.F., French, K.R., 1987. Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *Journal of Business* 60, 55–73.

Ghoshray, A., Johnson, B., 2010. Trends in world energy prices. *Energy Economics* 32, 1147–1156.

Gorton, G.B., Hayashi, F., Rouwenhorst, K.G., 2013. The fundamentals of commodity futures returns. *Review of Finance* 17, 35–105.

Gospodinov, N., Ng, S., 2013. Commodity prices, convenience yields, and inflation. *Review of Economics and Statistics* 95, 206–219.

Gouriéroux, C., Jasiak, J., 2016. Semi-parametric estimation of noncausal vector autoregression. CREST, unpublished manuscript.

Gouriéroux, C., Zakoian, J.M., 2016. Local explosion modelling by non-causal process. *Journal of Royal Statistical Society, Series B*, forthcoming.

Gronwald, M., 2016. Explosive oil prices. *Energy Economics* 60, 1–5.

Gutierrez, L., 2013. Speculative bubbles in agricultural commodity markets. *European Review of Agricultural Economics* 40, 217–238.

Hansen, L.P., Sargent, T.J., 1991. Two difficulties in interpreting vector autoregressions. In L.P. Hansen, T.J. Sargent (eds), *Rational Expectations Econometrics*, Westview Press, Boulder, 77–120.

Hencic, A., Gouriéroux, C., 2015. Noncausal autoregressive model in application to Bitcoin/USD exchange rates. In V.-N. Huynh et al. (eds.), *Econometrics of Risk, Studies in Computational Intelligence*, 17–40.

Karapanagiotidis, P., 2014. Dynamic modeling of commodity futures prices. MPRA

working paper No. 56805.

Kasa, K., Walker, T.B., Whiteman, C.H., 2014. Heterogeneous beliefs and tests of present value models. *Review of Economic Studies* 81, 1137–1163.

Kellard, N., Wohar, M.E., 2006. On the prevalence of trends in primary commodity prices. *Journal of Development Economics* 79, 146–167.

Lanne, M., Nyberg, H., 2015. Nonlinear dynamic interrelationships between real activity and stock returns. CREATES Research Paper 2015-36.

Lanne, M., Saikkonen, P., 2011. Noncausal autoregressions for economic time series. *Journal of Time Series Econometrics* 3, article 2.

Lanne, M., Saikkonen, P., 2013. Noncausal vector autoregression. *Econometric Theory* 29, 447–481.

Lanne, M., Luoto, J., Saikkonen, P., 2012a. Optimal forecasting of noncausal autoregressive time series. *International Journal of Forecasting* 28, 623–631.

Lanne, M., Nyberg, H., Saarinen, E., 2012b. Does noncausality help in forecasting economic time series? *Economics Bulletin* 32, 2849–2859.

Lof, M., 2013. Noncausality and asset pricing. *Studies in Nonlinear Dynamics and Econometrics* 17, 211–220.

Lof, M., 2014. GMM estimation with noncausal instruments under rational expectations. *Oxford Bulletin of Economics and Statistics* 76, 279–286.

Mamatzakis, E., Remoundos, P., 2011. Testing for adjustment costs and regime shifts in BRENT crude futures market. *Economic Modelling* 28, 1000–1008.

Nyberg, H., Saikkonen, P., 2014. Forecasting with a noncausal VAR model. *Computational Statistics and Data Analysis* 76, 536–555.

Pindyck, R.S., Rotemberg, J.J., 1990. The excess co-movement of commodity prices. *Economic Journal* 100, 1173–1189.

Taylor, M.P., Peel, D.A., Sarno, L., 2001. Nonlinear mean-reversion in real exchange rates: Toward a solution to the Purchasing Power Parity puzzles. *International Economic*

Review 42, 1015–1042.

Teräsvirta, T., 1994. Specification, estimation, and evaluation of smooth transition autoregressive models. *Journal of the American Statistical Association*, 89, 208–218.

Timmermann, A., 2006. Forecast combinations. In G. Elliott, C. Granger and A. Timmermann (eds.), *Handbook of Economic Forecasting* 1, 135–196.

Wang, Y., Wua, C., Yang, L., 2014. Oil price shocks and agricultural commodity prices. *Energy Economics* 44, 22–35.

West, K.D., 1996. Asymptotic inference about predictive ability. *Econometrica* 64, 1067–1084.

West, K.D., Wong, K.F., 2014. A factor model for co-movements of commodity prices. *Journal of International Money and Finance* 42, 289–309.

White, H., Granger, C.W.J., 2011. Consideration of trends in time series. *Journal of Time Series Econometrics* 3(1).