



Munich Personal RePEc Archive

Forecasting Base Metal Prices with Commodity Currencies

Pablo Pincheira and Nicolas Hardy

School of Business, Universidad Adolfo Ibáñez, School of Business,
Universidad Adolfo Ibáñez

2 January 2018

Online at <https://mpra.ub.uni-muenchen.de/83564/>

MPRA Paper No. 83564, posted 2 January 2018 23:06 UTC

Forecasting Base Metal Prices with Commodity Currencies[☆]

Pablo Pincheira Brown¹

Business School, Universidad Adolfo Ibáñez

Nicolás Hardy

Business School, Universidad Adolfo Ibáñez

January 2018

Abstract

In this paper we show that the Chilean exchange rate has the ability to predict the returns of the London Metal Exchange Index and of the six primary non-ferrous metals that are part of the index: aluminum, copper, lead, nickel, tin and zinc. The economic relationship hinges on the present-value theory for exchange rates, a floating exchange rate regime and the fact that copper represents about a half of Chilean exports and nearly 45% of Foreign Direct Investment. Consequently, the Chilean peso is heavily affected by fluctuations in the copper price. As all six base metal prices show an important comovement, we test whether the relationship between copper prices and Chilean exchange rates also holds true when it comes to the six primary non-ferrous metals. We find interesting evidence of predictability both in-sample and out-of-sample. Our paper is part of a growing literature that in the recent years has evaluated and called into question the ability of commodity currencies to forecast commodity prices.

Keywords:

Forecasting, commodities, prices, univariate time-series models, out-of-sample comparison, exchange rates, copper, primary non-ferrous metals

JEL: C52, C53, G17, E270, E370, F370, L740, O180, R310.

[☆]We thank Andrea Bentancor and Nikos Askitas for valuable comments. We also thank Michael McCracken for sharing tables of critical values. Any errors or omissions are responsibility of the authors.

*Corresponding author: Pablo Pincheira Brown.

¹Avenida Diagonal Las Torres 2640, Edificio C, Oficina 521-C, Peñalolén, Santiago. Chile. Tel.: +56 2 23311489.

1. Introduction

In this paper we show that the Chilean exchange rate has the ability to predict the returns of the London Metal Exchange Index (LMEX or LME Index) and of the six primary non-ferrous metals that are part of the index: aluminum, copper, lead, nickel, tin and zinc². Our results are important in two dimensions. First, they provide empirical evidence of a tool to forecast the commodity prices under analysis, and second, they are consistent with the present-value model for exchange rate determination.

Relatively recent literature has explored the empirical implications of the present-value model for exchange rates in commodity exporters countries. One of these key empirical implications is that exchange rates should have the ability to predict their fundamentals, see Engel and West (2005) and Campbell and Shiller (1987). In commodity exporter countries, one of these fundamentals is the price of the commodities being exported. This implies that “commodity currencies” should have the ability to predict the respective commodity price. This theoretical implication is extremely relevant given the first order impact of the fluctuations in commodity prices at the local and international level (see, for instance, Byrne, Fazio and Fiess, 2013; Arezki, Loungani, Van der Ploeg and Venables, 2014; and Groen and Pesenti, 2011) and the well known difficulty in forecasting asset returns and commodity prices, as it is well documented by Goyal and Welch (2008), Timmermann (2008), Reeve and Vigfusson (2011), Dooley and Lenihan (2005) and Engel and Valdés (2001) to name a few³.

Two interesting articles investigating this type of relationship are those of Chen, Rossi and Rogoff (2010, 2014) (hereafter CRR). In these papers the authors focus on five commodity exporting countries: Australia, Canada, Chile, New Zealand and South Africa. They show that these commodity currencies have the ability to forecast global commodity prices, and that some country specific exchange rates are able to forecast the price fluctuations of their own country’s commodity export bundle.

The results shown by CRR have received a lot of attention in the literature and, according to our review, researchers around the world have taken both sides of the debate, with supporters and detractors of the so called Commodity Currency Hypothesis (CCH). On the one hand, Chen, Rossi and Rogoff (2011) find more evidence of their previous findings in the specific case of agricultural commodity prices. Similarly, Gargano and Timmermann (2014) also show results supporting the CCH using the Australian dollar and the Indian rupee. Less enthusiastic is the paper by Groen and Pesenti (2011) with results that only mildly corroborate those of CRR when investigating the ability of commodity currencies to forecast ten alternative commodity indices. On the other hand, the findings of CRR have been called into question by Bork, Rovira and Sercu (2014) and more recently by Lof and Nyberg (2017).

²The LME index considers the following weights: aluminum (42.8%), copper (31.2%), zinc (14.8%), lead (8.2%), nickel (2%) and tin (1%).

³It is important to mention that while the papers cited here all coincide in that it is difficult to forecast commodity prices, there is another branch of the literature claiming important improvements in this task. See for instance the papers by Gargano and Timmermann (2014); Buncic and Moretto (2015); Liu, Hu, Li and Liu (2017); Sánchez, De Cos Juez, Suárez, Krzemień and Riesgo (2015); and He, Lu, Zou and Lai (2015).

These articles put forth different arguments and present empirical evidence, using basically the same data as in CRR, suggesting a relatively poor performance of commodity currencies when forecasting commodity prices.

Out of the five commodity currencies analyzed by CRR, the case of Chile is particularly interesting, as the bundle of commodity exports only includes copper. This makes sense as this mineral represents about a half of Chilean exports and nearly 45% of Foreign Direct Investment (FDI)⁴. Consequently, the Chilean peso is heavily affected by fluctuations in the price of this commodity. The implication of the present-value relationship is that the Chilean peso should have the ability to forecast the copper price. While Chen, Rossi and Rogoff (2010) provide relatively robust evidence supporting this view, in the update of their article released in 2014, they show much weaker out-of-sample results. As a matter of fact, they show that models including returns of the Chilean peso are only able to outperform the forecasts coming from the random walk with drift but are unsuccessful to outperform the simple AR(1) or the driftless random walk. In sharp contrast, for the specific case of the Chilean peso and the copper price, results in Bork, Rovira and Sercu (2014) are more compelling and provide evidence in favor of the CCH. These conflicting results call naturally for more research to analyze the potential predictive relationship between the Chilean peso and copper prices.

Copper is usually considered as belonging to a larger group of metals known as “base metals” or “non-ferrous metals” that, aside from copper, includes, aluminum, lead, nickel, tin and zinc. These metals have some similarities that are described in Roberts (2008, 2009) and Rossen (2015). More importantly, all these commodities are traded on a daily basis in the London Metal Exchange and they are also connected by the LME Index and the futures contracts linked to it. These linkages, plus the well known comovement in commodity prices reported for instance in West and Wong (2014), lead us to wonder whether the potential relationship between copper and the Chilean peso can also be generalized to the rest of the non ferrous metals and to the LME Index.

Our findings are fairly consistent with the major message in the first article by Chen, Rossi and Rogoff (2010). In sample and out-of sample tests provide evidence of predictability from the Chilean peso to all six base metals returns including those from the LME Index. Reductions in out-of-sample Mean Squared Prediction Errors tend to be small and unstable for some metals, however, although not for all of them. Notably, more stable and higher gains in forecast accuracy are achieved for the LME Index. While our methods follow closely those of the previous literature, there are also some differences that are worth mentioning. First, we focus on a sample period of floating Chilean exchange rates only, at both monthly and quarterly frequencies. Second, our data are converted to either monthly or quarterly frequencies by sampling from the last day of the month or the last day of the quarter. Finally, we use slightly different specifications than in the previous literature. We think that these details may be playing a role in explaining the robustness of our results.

⁴According to the Central Bank of Chile, copper exports represents 50.4% of total Chilean exports during the period January 2003-October 2017. Similarly, the mining sector has absorbed 44.3% of total FDI in Chile during the period 2009-2015.

The rest of this paper is organized as follows. In section 2 we present our data and forecasting models. In section 3 we present the results of our in-sample and out-of-sample exercises. Finally, in section 4 we present our conclusions.

2. Data and forecasting models

We consider monthly and quarterly data on Chilean exchange rates against the US dollar. Data in the same frequencies for all six commodity prices included in the LME Index are also considered. We include data for the LME Index as well. The source of our data is Thomson Reuters Datastream from which we downloaded the daily close price of each asset. Our data are converted to either monthly or quarterly frequencies by sampling from the last day of the month or the last day of the quarter. Consequently our empirical exercises are not polluted by the possible autocorrelation induced by taking monthly or quarterly averages of daily values, as mentioned in Bork, Rovira and Sercu (2014).

Our sample period goes from October 1999 until June 2017 at the monthly frequency, and from 1999Q4 until 2017Q2 at the quarterly frequency. This gives a total of 213 monthly observations and 71 quarterly observations⁵. The starting point of our sample period is determined by the date in which monetary authorities in Chile decided to pursue a pure flotation exchange rate regime. By the end of the nineties, Chile had achieved stable inflation and authorities were taking quick steps toward the implementation of a full fledged inflation targeting regime. See for instance the interesting description of this process in Morandé and Tapia (2002). Since then, the Chilean exchange rate has floated with only a few specific periods of pre-announced interventions that are described in Pincheira (2013a). Consequently, most of our sample period considers pure flotation of the Chilean peso. Some descriptive statistics of our series are found in Table A1 in the appendix.

Our econometric specifications are quite simple. They are inspired in the benchmarks used by CRR and by a vast literature that has shown that either the Random Walk (henceforth RW) or simple autoregressions are usually difficult benchmarks to beat when forecasting asset returns⁶. Our in-sample and out-of-sample analyses at the monthly frequency are based on the following simple specifications:

⁵In terms of one period returns, we have 212 monthly observations and 70 quarterly observations.

⁶Examples in the exchange rate literature are given by Meese and Rogoff (1983) and Clark and West (2006) for instance. When forecasting commodity prices, Chenn, Rossi and Rogoff (2010, 2014) consider the RW and an AR(1) as benchmarks, Lof and Nyberg (2017) consider both a causal and noncausal AR(1) whereas Groen and Pesenti (2011) uses autoregressions with more lags, but autoregressions in the end. Also Buncic and Moretto (2015) consider the RW as the benchmark to beat when forecasting copper prices. In stock returns, Goyal and Welch (2008) use the historical average as the benchmark when predicting excess returns.

Table 1: Monthly Specifications

1:	$\Delta \ln(CP_t) = c + \beta_1 \Delta \ln(ER_{t-1}) + \beta_2 \Delta \ln(ER_{t-2}) + \beta_{12} \Delta \ln(ER_{t-12}) + \rho \Delta \ln(CP_{t-1}) + \varepsilon_t$
2:	$\Delta \ln(CP_t) = c + \beta_1 \Delta \ln(ER_{t-1}) + \beta_2 \Delta \ln(ER_{t-2}) + \beta_{12} \Delta \ln(ER_{t-12}) + \varepsilon_t$
3:	$\Delta \ln(CP_t) = \beta_1 \Delta \ln(ER_{t-1}) + \beta_2 \Delta \ln(ER_{t-2}) + \beta_{12} \Delta \ln(ER_{t-12}) + \varepsilon_t$

Source: Authors' elaboration

where

$$\begin{aligned} \Delta \ln(CP_t) &\equiv \ln(CP_t) - \ln(CP_{t-1}) \\ \Delta \ln(ER_t) &\equiv \ln(ER_t) - \ln(ER_{t-1}) \end{aligned}$$

CP_t stands for ‘‘Commodity Price’’ and represents the generic predictand at time t , which in our case represents aluminum, copper, lead, nickel, tin, zinc and the LME Index. Similarly, ER_t represents the Chilean peso exchange rate at time t . This variable gives the amount of Chilean pesos needed to buy an american dollar in the domestic market. Finally, ε_t represents an error term.

Our quarterly specifications are slightly different and they are presented in Table 2.

Table 2: Quarterly Specifications

4:	$\Delta \ln(CP_t) = c + \beta [\Delta \ln(ER_{t-1}) + \Delta \ln(ER_{t-2})] + \rho \Delta \ln(CP_{t-1}) + \varepsilon_t$
5:	$\Delta \ln(CP_t) = c + \beta [\Delta \ln(ER_{t-1}) + \Delta \ln(ER_{t-2})] + \varepsilon_t$
6:	$\Delta \ln(CP_t) = \beta [\Delta \ln(ER_{t-1}) + \Delta \ln(ER_{t-2})] + \varepsilon_t$

Source: Authors' elaboration

We notice here that in the quarterly specifications we have imposed the restriction that the coefficients associated to the first and second lags of Chilean exchange rate returns are equal. We do that for two reasons. First, our exploratory in-sample and out-of sample exercises suggest that this is a good idea⁷. Second, we are looking for parsimony, given the relatively low number of quarterly observations.

For these specifications, we consider the following null hypotheses H_{01} and H_{02} at the monthly and quarterly frequency respectively:

$$\begin{aligned} H_{01} &: \beta_1 = \beta_2 = \beta_{12} = 0 \\ H_{02} &: \beta = 0 \end{aligned}$$

⁷In sample evaluations indicate that the coefficients associated to the first two lags are negative and, in many cases, the null hypotheses that both coefficients are equal cannot be rejected at usual significance levels.

This means that we are comparing our models to the simple AR(1) for $\Delta \ln(CP_t)$ and to the RW with and without drift for $\ln(CP_t)$. Our null hypotheses H_{01} and H_{02} posit that the Chilean exchange rate has no role in predicting commodity returns. We test these hypotheses both in-sample and out-of sample focusing on one-step-ahead forecasts only, leaving the analysis of multistep ahead forecasts as an extension for future research.

In sample evaluations are simply carried out using a Wald statistic for monthly regressions, given that we have three lags of Chilean exchange rate returns. At the quarterly frequency we consider the simple t-statistic associated to the coefficient of Chilean exchange rate returns.

For out-of-sample evaluations we use the ENCNEW test proposed by Clark and McCracken (2001) and also used by CRR. This test has a non-standard asymptotic distribution, but critical values for one-step-ahead forecasts are available in the appendix of that paper. In general, Clark and McCracken (2001) show that the resulting asymptotic distribution of the ENCNEW test is a functional of Brownian motions depending on the number of excess parameters of the nesting model, which is 1 in our quarterly regressions and 3 in our monthly regressions, the parameter π defined as the limit of the ratio P/R , where P is the number of one-step-ahead forecasts and R is the size of the first expanding window used in the out-of-sample analyses. The asymptotic distribution of the test varies also with the scheme used to update the estimates of the parameters: either rolling, expanding or fixed. See Clark and McCracken (2001) or West (2006) for further details about the implementation of out-of-sample tests of predictive ability in nested environments.

For our in-sample analysis we estimate our models with all the available observations. For the out-of-sample analysis we split the sample in two windows: an initial estimation window of size R and a prediction window of size P such that $P + R = T$, where T is the total number of observations. We consider two different ways of splitting our samples. First we use approximately one third of our observations for initial estimation and two thirds for evaluation. This means that we pick $R = 71$ in our monthly analysis and $R = 23$ in our quarterly analysis. As a robustness check we also consider a different situation in which we pick approximately 70% of our sample for initial estimation and 30% for evaluation. This means $R = 152$ in our monthly analysis and $R = 51$ in our quarterly analysis⁸. We also carry out inference using the ENC-t test of Clark and McCracken (2001), which is based on the same statistic analyzed in Clark and West (2007) and used in Groen and Pesenti (2011). Results are slightly less powerful than those obtained with the ENCNEW test, which is consistent with the simulation evidence in Clark and McCracken (2001, 2005) showing that the ENCNEW test tends to have more power than the ENC-t test. As a consequence, we do not show tables with the ENC-t test, but they are available upon request.

⁸The short estimation window spans the period October 1999-August 2005 at the monthly frequency and the period 1999Q4-2005Q2 at the quarterly frequency. The longer estimation window spans the period October 1999-July 2011 at the monthly frequency and the period 1999Q4-2012Q2 at the quarterly frequency. These choices are such that the outburst of the financial crisis in 2008 is well in the middle of our estimation or prediction windows. Furthermore, this choice is partly driven by the table of asymptotic critical values in Clark and McCracken (2001). This table provides critical values for specific values of the ratio P/R . We consider two values: 2 and 0.4.

3. Empirical Results

In this section we report in sample estimates and tests of equations 1 and 4 in Tables 1 and 2. We also show results of the ENCNEW out-of-sample test of Clark and McCracken (2001). We start by reporting our in-sample results.

3.1. In sample analysis

Tables 3 and 4 next show the estimates of equations 1 and 4 in Tables 1 and 2. In Table 3 we consider monthly frequencies, whereas in Table 4 we focus on quarterly frequencies. In both Tables 3 and 4 we use HAC standard errors according to Newey and West (1987, 1994).

Table 3: Forecasting Base Metals Returns with the Chilean Peso
In-sample analysis with monthly data

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	<i>Aluminum</i>	<i>Copper</i>	<i>Lead</i>	<i>Nickel</i>	<i>Tin</i>	<i>Zinc</i>	<i>Lmex</i>
<i>CLP(-1)</i>	-0.407*** (0.146)	-0.419** (0.203)	-0.787*** (0.276)	-0.458* (0.250)	-0.304** (0.158)	-0.533*** (0.225)	-0.403*** (0.162)
<i>CLP(-2)</i>	-0.270* (0.159)	-0.456*** (0.175)	-0.241 (0.218)	0.0149 (0.214)	-0.470*** (0.171)	-0.137 (0.153)	-0.323*** (0.145)
<i>CLP(-12)</i>	0.359*** (0.124)	0.312** (0.150)	0.557*** (0.174)	0.467** (0.212)	0.409** (0.181)	0.542*** (0.160)	0.367** (0.123)
<i>Aluminum(-1)</i>	-0.066 (0.087)						
<i>Copper(-1)</i>		0.054 (0.095)					
<i>Lead(-1)</i>			-0.122 (0.080)				
<i>Nickel(-1)</i>				0.001 (0.079)			
<i>Tin(-1)</i>					0.048 (0.081)		
<i>Zinc(-1)</i>						-0.094 (0.082)	
<i>Lmex(-1)</i>							0.017 (0.092)
<i>Constant</i>	0.002 (0.004)	0.006 (0.006)	0.009 (0.006)	0.001 (0.007)	0.007 (0.005)	0.005 (0.006)	0.004 (0.005)
Observations	200	200	200	200	200	200	200
R-squared	0.096	0.099	0.103	0.04	0.100	0.082	0.105
Wald test p-value	0.001	0.003	0.001	0.058	0.002	0.003	0.001

CLP stands for Chilean Peso Returns

Table 3 presents estimations of equation 1 in Table 1. * p<10%, ** p<5%, *** p<1%

Source: Authors' elaboration.

In the last row of Table 3 we report the p-value of the Wald statistic associated to the null hypotheses that all the three coefficients of Chilean peso returns are zero. All entries indicate that this null hypothesis is rejected at the 10% significance level. Moreover, leaving aside the case of Nickel, these null hypotheses are rejected at the 1% significant level.

Several additional features are worth noticing. First, neither the constant term nor the first lag of the asset returns being predicted are statistically significant. This is consistent with very little autocorrelation in asset returns. Second, the first lag of Chilean peso returns is negative and statistically significant in all assets. In contrast, the twelfth lag is positive in all entries and statistically significant as well. With the second lag we find mixed results in terms of statistical significance although most of them are negative. Notably, the coefficients associated to the Chilean peso are, in general, high in absolute value. Finally, the coefficient of determinations are in the range of 4%-11%, being higher for lead, copper and the LME Index, and lower in the case of Tin. All in all, our monthly results are very impressive and suggest an interesting ability of the Chilean peso to predict base metals returns.

Table 4: Forecasting Base Metals Returns with the Chilean Peso
In-sample analysis with quarterly data

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	<i>Aluminum</i>	<i>Copper</i>	<i>Lead</i>	<i>Nickel</i>	<i>Tin</i>	<i>Zinc</i>	<i>Lmex</i>
<i>CLP(-1) + CLP(-2)</i>	-0.476** (0.211)	-0.771** (0.380)	-0.805*** (0.257)	-0.919*** (0.280)	-0.907*** (0.223)	-0.699*** (0.273)	-0.917*** (0.250)
<i>Aluminum(-1)</i>	0.027 (0.134)						
<i>Copper(-1)</i>		-0.132 (0.221)					
<i>Lead(-1)</i>			-0.084 (0.141)				
<i>Nickel(-1)</i>				-0.077 (0.120)			
<i>Tin(-1)</i>					-0.001 (0.123)		
<i>Zinc(-1)</i>						0.029 (0.165)	
<i>Lmex(-1)</i>							-0.120 (0.120)
<i>Constant</i>	0.006 (0.012)	0.025 (0.020)	0.033* (0.019)	0.006 (0.023)	0.026 (0.017)	0.017 (0.017)	0.019 (0.013)
Observations	68	68	68	68	68	68	68
R-squared	0.152	0.108	0.154	0.134	0.269	0.158	0.285

CLP stands for Chilean Peso Returns

Table 4 presents estimations of equation 4 in Table 2. * p<10%, ** p<5%, *** p<1%

Source: Authors' elaboration.

Results from Table 4 are similar and equally striking as those from Table 3. In particular, we find no signs of autocorrelation, and all the coefficients associated to Chilean peso returns are

statistically significant at tight significance levels. They are also all negative and important in magnitude. Now the coefficients of determinaton are higher, in the range of 10%-29%. Remarkably, the lowest coefficient is achieved for copper, and the highest for the LME Index. Of course this last point must not be considered too seriously as it is a textbook fact that the addition of irrelevant variables may induce an increment in the R^2 diagnostic statistic.

All in all, our in-sample estimates provide evidence of a relationship between the time series on base metal returns and the returns of the Chilean peso both at the monthly and quarterly frequency. In-sample estimates, however, are usually criticized because they are relatively different from a real time forecasting exercise and also because they are prone to overfitting. To mitigate these shortcomings, we move next to an out-of-sample analysis.

3.2. Out-of-sample Analysis

Tables 5-8 show results of the ENCNEW test of Clark and McCracken (2001) in different out-of-sample exercises. Tables 5 and 6 show results when the number of forecasts is twice the number of observations of the first estimation window ($\pi \equiv P/R = 2$). Tables 7 and 8 are similar to Tables 5 and 6 but they report results when the number of forecasts is only 40% of the number of observations used in the first estimation window ($\pi \equiv P/R = 0.4$). We consider expanding windows with monthly and quarterly data and, as mentioned before, two different sizes of the first estimation window.

Table 5: Forecasting Base Metals Returns with Chilean Exchange Rates
Out-of-sample analysis with monthly data, ($\pi \equiv P/R = 2$)

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	ENCNEW						
Benchmark Model	<i>aluminum</i>	<i>copper</i>	<i>lead</i>	<i>nickel</i>	<i>tin</i>	<i>zinc</i>	<i>lmex</i>
AR(1)	8.87***	5.42**	10.44***	5.32***	10.24***	6.51***	8.16***
RW	10.44***	12.34***	9.75***	7.29***	12.29***	8.08***	13.89***
DRW	10.64***	12.58***	9.93***	7.49***	12.40***	8.12***	14.13***

10%, 5% and 1% critical values are 2.336, 3.564 and 5.805 respectively for ENCNEW when excess parameters are 3.

P represents the number of one-step-ahead forecasts, R the sample size of the first estimation window.

The AR(1) benchmark corresponds to model 1 in Table 1 when all the coefficients associated to the Chilean peso are set to zero. Similarly, the RW and DRW benchmarks correspond to models 2 and 3 in Table 1, respectively, when all the coefficients associated to the Chilean peso are set to zero

* p<10%, ** p<5%, *** p<1%

Source: Authors' elaboration.

Table 6: Forecasting Base Metals Returns with Chilean Exchange Rates
Out-of-sample analysis with quarterly data, ($\pi \equiv P/R = 2$)

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ENCNEW							
Benchmark Model	<i>aluminum</i>	<i>copper</i>	<i>lead</i>	<i>nickel</i>	<i>tin</i>	<i>zinc</i>	<i>lmex</i>
AR(1)	3.84**	4.79***	5.11***	3.94**	12.14***	3.74**	7.27***
RW	6.28***	2.72**	5.96***	5.37***	15.10***	5.86***	12.35***
DRW	6.31***	2.78**	5.89***	5.46***	14.34***	5.70***	12.22***

10%, 5% and 1% critical values are 1.280, 2.085 and 4.134 respectively for ENCNEW when excess parameters are 1.

P represents the number of one-step-ahead forecasts, R the sample size of the first estimation window.

The AR(1) benchmark corresponds to model 4 in Table 2 when the coefficient associated to the Chilean peso is set to zero. Similarly, the RW and DRW benchmarks correspond to models 5 and 6 in Table 2,

respectively, when the coefficient associated to the Chilean peso is set to zero

* p<10%, ** p<5%, *** p<1%

Source: Authors' elaboration.

The null hypothesis of no predictability is rejected in all cases at least at the 10% significant level. This means that we detect predictability against all our three benchmarks models, for all six base metals and also for the LME Index. The evidence of predictability shown in Tables 5-8 is also robust to the frequency of our data (either monthly or quarterly) and to the choice of the point in time in which we split our sample. Moreover, in many entries of the tables, we reject the null hypothesis at tight significance levels of 1%.

Table 7: Forecasting Base Metals Returns with Chilean Exchange Rates
Out-of-sample analysis with monthly data, ($\pi \equiv P/R = 0.4$)

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ENCNEW							
Benchmark Model	<i>aluminum</i>	<i>copper</i>	<i>lead</i>	<i>nickel</i>	<i>tin</i>	<i>zinc</i>	<i>lmex</i>
AR(1)	9.74***	5.82***	11.97***	4.74***	8.66***	8.01***	10.01***
RW	6.51***	4.02***	7.42***	4.09***	7.82***	6.09***	7.53***
DRW	6.47***	3.68***	7.31***	4.23***	7.72***	6.04***	7.45***

10%, 5% and 1% critical values are 1.285, 1.865 and 3.098 respectively for ENCNEW when excess parameters are 3.

P represents the number of one-step-ahead forecasts, R the sample size of the first estimation window.

The AR(1) benchmark corresponds to model 1 in Table 1 when all the coefficients associated to the Chilean peso are set to zero. Similarly, the RW and DRW benchmarks correspond to models 2 and 3 in Table 1,

respectively, when all the coefficients associated to the Chilean peso are set to zero

* p<10%, ** p<5%, *** p<1%

Source: Authors' elaboration.

Table 8: Forecasting Base Metals Returns with Chilean Exchange Rates
 Out-of-sample analysis with quarterly data, ($\pi \equiv P/R = 0.4$)

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ENCNEW							
Benchmark Model	<i>aluminum</i>	<i>copper</i>	<i>lead</i>	<i>nickel</i>	<i>tin</i>	<i>zinc</i>	<i>lmex</i>
AR(1)	2.41***	2.47***	2.70***	2.54***	3.13***	2.25***	6.22***
RW	1.52**	2.99***	1.41***	2.83***	2.73***	2.19***	6.83***
DRW	1.34**	2.33***	0.77*	2.54***	2.06**	1.66**	5.69***

10%, 5% and 1% critical values are 0.685, 1.079 and 2.098 respectively for ENCNEW when excess parameters are 1.

P represents the number of one-step-ahead forecasts, R the sample size of the first estimation window.

The AR(1) benchmark corresponds to model 4 in Table 2 when the coefficient associated to the Chilean peso is set to zero. Similarly, the RW and DRW benchmarks correspond to models 5 and 6 in Table 2,

respectively, when the coefficient associated to the Chilean peso is set to zero

* $p < 10\%$, ** $p < 5\%$, *** $p < 1\%$

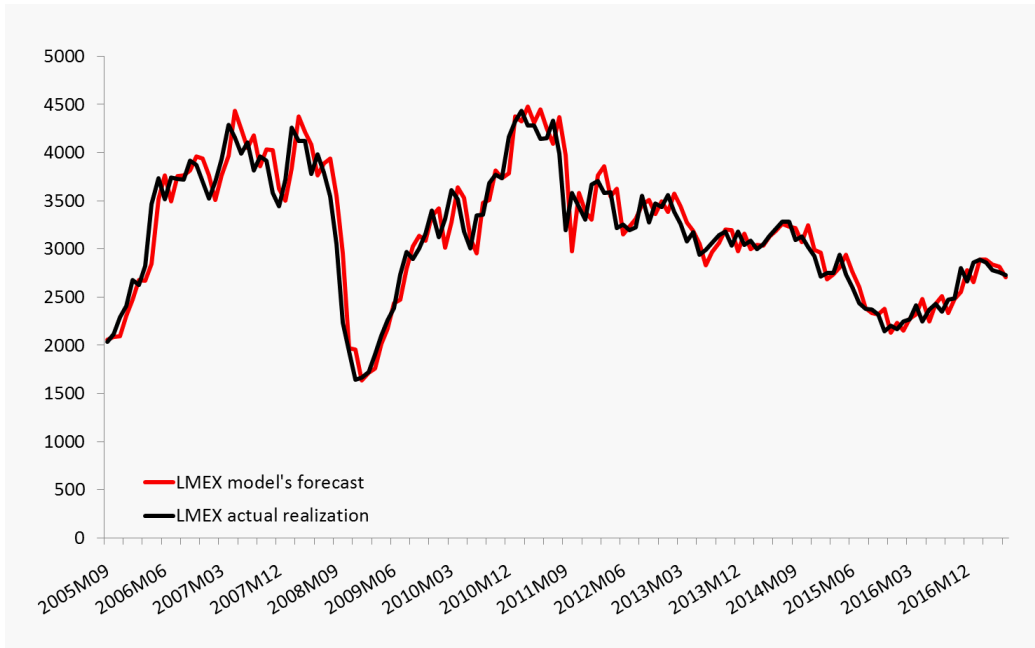
Source: Authors' elaboration.

Similar tables constructed when the out-of-sample analyses are carried out in rolling windows basically report the same qualitative results. They are not shown to save space, but they are available upon request.

We also carry out inference of predictive ability with the ENC-t test of Clark and McCracken (2001) which uses the same core statistic as the normal tests in Clark and West (2007, 2006). Results with this test are slightly less compelling than those obtained with the ENCNEW test. In fact, similar tables constructed with the ENC-t test reveal that the null hypothesis of no predictability is rejected in all cases but three: the case of nickel when the benchmark is the AR(1) and $P/R = 2$, and the case of aluminum when the benchmarks are both random walks with and without drift and $P/R = 0.4$.

To give an illustration of the accuracy of our approach, figures 1 and 2 show the forecasts constructed with model 2 in Table 1 when predicting the LME Index. The pictures also show the real historical values of the index for comparison. Figure 1 shows the forecasts in levels, while Figure 2 shows the forecasts of one month returns.

Picture 1: Forecasting the LME Index with Chilean Exchange Rates
Out-of-sample analysis with monthly data




One step ahead forecasts based on model 2 in Table 1. Sample period: May 2009 - June 2017

Source: Authors' elaboration.

Figure 1 is impressive, as our forecasts closely follows the LME index. We notice, however, that this picture is impressive in part because the level of the LME Index is implicitly predicted by a random walk plus Chilean exchange returns. To isolate the forecasting ability of the Chilean peso, Figure 2 shows the forecasts of the LME Index returns. This picture is less impressive but it shows a positive correlation between the forecasts generated with model 2 in Table 1 and actual LME Index returns. In fact the correlation is 0.25 suggesting an interesting predictive ability.

Picture 2: Forecasting LME Index Returns with Chilean Exchange Rates
Out-of-sample analysis with monthly data



One step ahead forecasts based on model 2 in Table 1. Sample period: May 2009 - June 2017

Source: Authors' elaboration.

Why our results are stronger than those of some previous literature like the update of CRR? First, it is important to emphasize that most of our results are new, as we are not aware of any other paper analyzing the predictive ability of the Chilean peso to all six base metals returns. Nevertheless, as mentioned before, our results seem stronger than those reported by CRR (2014). Several details may be behind these different results: first we use a sample period of pure flotation of the Chilean peso, whereas CRR (2010, 2014) consider 10 years of data in which the Chilean peso underwent constant interventions by the Central Bank. Second, our data are based on daily observations and not monthly or quarterly averages, and third, our specifications are slightly different in a relevant manner. For instance, in our monthly regressions we detect the relevance of the first, second and twelve lags of Chilean peso returns. If we would have only considered the first lag as CRR do, our results would have been very different. Table 9 next show the results with only one lag as in CRR (2010, 2014). In this case, it is only for nickel and lead that the models augmented with the Chilean exchange rate are able to outperform all three benchmarks. The case of copper is particularly interesting, as it is the only metal for which no rejection is found whatsoever. Results in Table 9 are obtained with expanding windows, but the use of rolling windows show the same picture, which is similar to that described in CRR (2014). These results show that our specifications seem to be more adequate and responsible in part for the differences between our results and those of the previous literature.

Table 9: Forecasting Base Metals Returns with One Lag of Chilean Exchange Rates Returns
 Out-of-sample analysis with monthly data, ($\pi \equiv P/R = 0.4$)

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ENCNEW							
Benchmark Model	<i>aluminum</i>	<i>copper</i>	<i>lead</i>	<i>nickel</i>	<i>tin</i>	<i>zinc</i>	<i>lmex</i>
AR(1)	2.06**	0.46	3.92***	1.45**	0.84*	1.06*	1.41**
RW	0.56	-0.11	1.83**	1.01*	0.77*	-0.01	0.63
DRW	0.49	-0.52	1.50**	0.95*	0.57	-0.13	0.32

10%, 5% and 1% critical values are 0.685, 1.079 and 2.098 respectively for ENCNEW when excess parameters are 1.

P represents the number of one-step-ahead forecasts, R the sample size of the first estimation window.

The AR(1) benchmark corresponds to model 1 in Table 1 when all the coefficients associated to the Chilean peso are set to zero. Similarly, the RW and DRW benchmarks correspond to models 2 and 3 in Table 1, respectively, when all the coefficients associated to the Chilean peso are set to zero

* p<10%, ** p<5%, *** p<1%

Source: Authors' elaboration.

3.3. Forecast Accuracy

Traditional out-of-sample tests of predictive ability in nested environments, like the ENC-NEW and ENC-t tests, are designed to compare the population Mean Squared Prediction Errors (MSPE) of two models under evaluation. Due to sampling error, however, the model displaying lower MSPE in population, may not necessarily display the lowest MSPE in a given sample. This is the reason behind the adjustment proposed in the tests developed by Clark and West (2006, 2007). Nevertheless, whenever the nested model is rejected by one of these tests, the implication is that a linear combination of the forecasts coming from the nesting and nested models should display the lowest sample MSPE (see Pincheira, 2013b, for further details in the particular case in which the nested model is the simple random walk). With this in mind, Tables 10 and 11 show out-of-sample coefficients of determination (R^2) according to the definition used by Goyal and Welch (2008). In other words, Tables 10 and 11 provide a measure of forecast accuracy of the model including exchange rate returns relative to the model that simply predicts commodity returns with a constant. We use models 2 and 4 from Tables 1 and 2 respectively to construct Tables 10 and 11. A number greater than zero means that the model using exchange rate returns is more accurate than the simple benchmark used by Goyal and Welch (2008). A zero entry indicates that both models produce similarly accurate forecasts. A negative figure indicates that the simple benchmark outperforms the model using Chilean exchange rate returns. The first row in Tables 10 and 11 shows in-sample R^2 for comparison.

Table 10: Out of Sample (OOS) R^2 when Forecasting Base Metals Returns with the Chilean Peso
 Out-of-sample analysis with monthly data

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	<i>Aluminum</i>	<i>Copper</i>	<i>Lead</i>	<i>Nickel</i>	<i>Tin</i>	<i>Zinc</i>	<i>Lmex</i>
<i>In Sample R2</i>	0.09	0.10	0.09	0.04	0.10	0.08	0.11
<i>OOS-Rolling, P/R=2</i>	0.03	0.01	0.00	-0.02	0.01	0.00	0.02
<i>OOS-Recursive, P/R=2</i>	0.03	0.02	0.01	0.00	0.03	0.01	0.03
<i>OOS-Rolling, P/R=0.4</i>	0.04	-0.02	0.04	0.03	0.06	0.02	0.02
<i>OOS-Recursive, P/R=0.4</i>	0.05	-0.02	0.04	0.03	0.06	0.03	0.02

P represents the number of one-step-ahead forecasts, R the sample size of the first estimation window.

Source: Authors' elaboration.

Table 11: Out of Sample (OOS) R^2 when Forecasting Base Metals Returns with the Chilean Peso
 Out-of-sample analysis with quarterly data

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	<i>Aluminum</i>	<i>Copper</i>	<i>Lead</i>	<i>Nickel</i>	<i>Tin</i>	<i>Zinc</i>	<i>Lmex</i>
<i>In Sample R2</i>	0.15	0.10	0.15	0.13	0.27	0.16	0.28
<i>OOS-Rolling, P/R=2</i>	0.02	-0.06	0.06	0.03	0.12	0.04	0.09
<i>OOS-Recursive, P/R=2</i>	0.05	-0.01	0.06	0.06	0.14	0.05	0.12
<i>OOS-Rolling, P/R=0.4</i>	-0.05	0.06	-0.01	0.06	0.05	0.04	0.12
<i>OOS-Recursive, P/R=0.4</i>	-0.04	0.03	0.00	0.06	0.04	0.04	0.12

P represents the number of one-step-ahead forecasts, R the sample size of the first estimation window.

Source: Authors' elaboration.

A quick inspection of Tables 10 and 11 reveals several interesting features. First, out-of-sample R^2 are lower than their in-sample counterparts. Second, while most of the entries display positive numbers, a few cells contain negative numbers, suggesting a poor performance of the models with Chilean exchange rate returns at the sample level. Third, out-of-sample R^2 tend to be low, with an average of 2% and 5% with monthly and quarterly data respectively. Fourth, there is some instability of out-of-sample R^2 across different out-of-sample exercises, with change in signs in a few cases. Notably, relatively robust results are detected in the cases of Tin, Zinc and the LME Index. All in all, results in Tables 10 and 11 indicate that the strong evidence of population predictive ability detected with the ENCNEW test, turns into quantitatively modest predictive ability at the sample level.

4. Concluding remarks

In recent years, and starting with the paper by Chen, Rossi and Rogoff (2010), a growing literature has evaluated and called into question the ability of commodity currencies to forecast commodity prices. In this paper, we contribute to this literature by showing that

the Chilean exchange rate has the ability to predict the returns of the London Metal Exchange Index and of the six primary non-ferrous metals that are part of the index: aluminum, copper, lead, nickel, tin and zinc. Using both in-sample and out-of-sample analysis we find strong evidence of this predictability at the population level, yet modest quantitative evidence at the sample level. These findings illustrate the need of new estimation and forecasting methods aimed at closing the existing gap between population and sample predictability. While the econometric literature has been successful in developing new tests of predictive ability in nested environments, it is still not entirely clear how a sound detection of predictability at the population level may be used to construct forecasts reflecting this predictability at the sample level.

Our results are important and timely, as a few relatively recent papers have reported results supporting the idea that exchange rates do not have the ability to forecast commodity prices. For instance, in the particular case of the literature exploiting the relationship between copper prices and the Chilean peso, the update of Chen, Rossi and Rogoff (2014) show weak results. In contrast, Bork, Rovira and Sercu (2014) show better results for this particular pair of commodity and exchange rate, but very poor results for other combinations. Lof and Nyberg (2017) point out in the same direction. While we do not expect to settle the debate, we hope to provide sound empirical evidence that can be used to enrich the discussion.

It is relatively easy to explain the ability that the Chilean peso has to predict copper prices in light of the present-value model for exchange rate determination. Yet, it is not that easy to sketch an explanation for the ability to predict all six base metal returns. In particular, a direct application of the same argument would not find empirical support, as the relevance of base metals in Chile's trade balance, other than copper, is negligible. A case can be made for aluminum, as it is a close substitute of copper in several industrial applications. Other than that, we would need to rely on the well known comovement in commodity prices or maybe on a dollar effect to find a plausible explanation (see Akram, 2009 and Chen, Jackson, Kim and Resiandini 2014). An interesting avenue for future research could address this point and offer more clarity about the predictive ability linking seemingly unrelated asset returns.

References

1. Akram, Q. (2009), Commodity prices, interest rates and the dollar. *Energy Economics*, Volume 31, Issue 6, November 2009, Pages 838-851.
2. Arezki, R, P. Loungani, R. Van der Ploeg and A. Venables (2014). Understanding international commodity price fluctuations. *Journal of International Money and Finance*. 42, 1-8.
3. Bork, L., P. Rovira and P. Sercu (2014). Do Exchange Rates Really Help Forecasting Commodity Prices? (August 20, 2014). Available at SSRN: <https://ssrn.com/abstract=2473624> or <http://dx.doi.org/10.2139/ssrn.2473624>
4. Buncic D. and C. Moretto (2015). Forecasting copper prices with dynamic averaging and selection models. *North American Journal of Economics and Finance* 33 (2015) 1-38.
5. Byrne J, G. Fazio and N. Fiess (2013). Primary commodity prices: Co-movements, common factors and fundamentals. *Journal of Developing Economies* 101 (2013) 16-26.

6. Campbell, J.Y. and R. J. Shiller (1987). Cointegration and Tests of Present Value Models, *Journal of Political Economy*, 1987, 95 (5), 1062–88.
7. Chen Y., K. Rogoff and B. Rossi (2010) Can Exchange Rates Forecast Commodity Prices?. *The Quarterly Journal of Economics*, August 2010, 125 (3), 1145–1194.
8. Chen Y., K. Rogoff and B. Rossi (2014) Can Exchange Rates Forecast Commodity Prices? An Update, manuscript, February 2014.
9. Chen Y., K. Rogoff and B. Rossi (2011) Predicting Agri-Commodity Prices: An Asset Pricing Approach, World Uncertainty and the Volatility of Commodity Markets, ed. B. Munier, IOS, 2011.
10. Chen S.L, J.D. Jackson, H. Kim and P. Resiandini (2014). What drives commodity prices? *American Journal of Agricultural Economics* 96, 1455–1468.
11. Clark T. E. and M. W. McCracken (2001). Tests of equal forecast accuracy and encompassing for nested models. *Journal of Econometrics*, 105(1), 85–110.
12. Clark T. E. and M. W. McCracken (2005). Evaluating Direct Multistep Forecasts. *Econometric Reviews*, 24(4), 369–404.
13. Clark T.E. and K.D. West (2006). Using Out-of-Sample Mean Squared Prediction Errors to Test the Martingale Difference Hypothesis. *Journal of Econometrics* 135: 155-186.
14. Clark T. E. and K. D. West (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, 138(1), 291–311
15. Dooley G. and H. Lenihan (2005). An assessment of time series methods in metal price forecasting. *Resources Policy* 30 (2005) 208–217.
16. He K, X. Lu, Y. Zou and K. Lai (2015). Forecasting metal prices with a curvelet based multiscale methodology. *Resources Policy* 45 (2015) 144–150.
17. Engel C. and K. D. West (2005). Exchange Rates and Fundamentals. *Journal of Political Economy*, June 2005, 113 (3), 485–517
18. Engel E. and R. Valdés (2001), Predicting copper prices: Beyond random walk? , CEA Working Paper N°100, Universidad de Chile (in spanish).
19. Gargano A. and A. Timmermann (2014) Forecasting Commodity Price Indexes Using Macroeconomic and Financial Predictors. *International Journal of Forecasting* 30, 825-843.
20. Goyal A. and I. Welch (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4), 1455–1508.
21. Groen J. and P. Pesenti. (2011). Commodity prices, commodity currencies, and global economic developments. In T. Ito, &A. K. Rose (Eds.), (Vol. 20) Commodity prices and markets, NBER – East Asia seminar on economics (pp. 15–42). University of Chicago Press.
22. Liu C, Z. Hu, Y. Li and S. Liu (2017). Forecasting copper prices by decision tree learning, *Resources Policy* 52 (2017) 427–434.
23. Lof M. and H. Nyberg (2017). Noncausality and the commodity currency hypothesis. *Energy Economics* 65 (2017) 424–433.
24. Morandé F. G. and M. Tapia (2002) “Exchange Rate Policy in Chile: From the Band to Floating and Beyond,” Working Papers Central Bank of Chile 152, Central Bank of Chile April 2002.

25. Meese R. and K. Rogoff. (1983). Empirical Exchange Rate Models of the Seventies. Do They Fit Out-of-Sample? *Journal of International Economics* 14: 3-24.
26. Newey, W. and K.D. West. (1987). A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3), 703-708.
27. Newey, W. and K.D. West. (1994) Automatic Lag Selection in Covariance Matrix Estimation. *The Review of Economic Studies*. 61(4): 631-53.
28. Pincheira P. (2013a). Interventions and inflation expectations in an inflation targeting economy. BIS Working Papers 427, Bank for International Settlements.
29. Pincheira P. (2013b). Shrinkage-Based Tests of Predictability, *Journal of Forecasting*, John Wiley & Sons, Ltd., vol. 32(4), pages 307-332, July.
30. Reeve T. and R. J. Vigfusson (2011). Evaluating the Forecasting Performance of Commodity Futures Prices. Board of Governors of the Federal Reserve System, International Finance Discussion Papers N°1025, August 2011.
31. Roberts, M. (2008). Synchronization and Co-movement of Metal Prices. *Minerals & Energy*, 23 (3), 105–118.
32. Roberts, M. (2009). Duration and Characteristics of Metal Price Cycles. *Resources Policy*, 34 (3), 87–102.
33. Rossen, A. (2015). What are metal prices like? Co-movement, price cycles and long-run trends, Hamburg Institute of International Economics Working Paper 155.
34. Sánchez F, F.De Cos Juez, A. Suárez, A. Krzemień and P. Riesgo (2015). Forecasting the COMEX copper spot price by means of neural networks and ARIMA models. *Resources Policy* 45 (2015) 37–43.
35. Timmermann, A. (2008). Elusive return predictability. *International Journal of Forecasting*, 24(1), 1-18.
36. West K.D. and K.F. Wong (2014). A factor model for co-movements of commodity prices. *Journal of International Money and Finance* 42 (2014) 289–309
37. West K.D. (2006). Forecast Evaluation. In Handbook of Economic Forecasting Volume 1. G. Elliot, C. Granger and A. Timmermann editors. Elsevier.

Appendix A. Descriptive Statistics

Table A1: Descriptive Statistics of the Monthly and Quarterly Returns of our Series

Sample Period								
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Monthly Data								
	<i>Aluminum</i>	<i>Copper</i>	<i>Lead</i>	<i>Nickel</i>	<i>Tin</i>	<i>Zinc</i>	<i>Lnex</i>	<i>Clp</i>
Mean	0.0011	0.0055	0.0069	0.0005	0.0058	0.0038	0.0035	0.0009
Median	0.0001	0.0071	0.0089	-0.0020	0.0061	0.0051	0.0057	0.0006
Standard								
Deviation	0.0594	0.0774	0.0903	0.1030	0.0725	0.0795	0.0632	0.0329
Max	0.1564	0.2709	0.2399	0.3009	0.2382	0.2450	0.2030	0.1988
Min	-0.1686	-0.4208	-0.3231	-0.2965	-0.2093	-0.3802	-0.3110	-0.0699
Quarterly Data								
	<i>Aluminum</i>	<i>Copper</i>	<i>Lead</i>	<i>Nickel</i>	<i>Tin</i>	<i>Zinc</i>	<i>Lnex</i>	<i>Clp</i>
Mean	0.0020	0.0159	0.0220	0.0026	0.0179	0.0120	0.0114	0.0032
Median	0.0072	0.0236	0.0202	-0.0056	0.0127	0.0092	0.0185	0.0080
Standard								
Deviation	0.1067	0.1637	0.1607	0.2000	0.1452	0.1512	0.1303	0.0597
Max	0.2168	0.3284	0.4663	0.4605	0.3582	0.4286	0.4270	0.1827
Min	-0.4867	-0.7957	-0.3585	-0.5297	0.3741	0.4756	-0.5538	-0.1304

Source: Authors' elaboration.