# The Mean Squared Prediction Error Paradox: A summary

Pincheira, Pablo and Hardy, Nicolas

Universidad Adolfo Ibañez, Universidad Finis Terrae

29 December 2020

# Summary of the paper entitled "The Mean Squared Prediction Error Paradox"[1]

Pablo Pincheira Brown[♣]

School of Business

Universidad Adolfo Ibáñez

Nicolás Hardy[§]

School of Economics and Business

Universidad Finis Terrae

**December 2020**

**Abstract**

This is a summary of the paper entitled[2]: "The Mean Squared Prediction Error Paradox". In that paper, we show that traditional comparisons of Mean Squared Prediction Error (MSPE) between two competing forecasts may be highly controversial. This is so because when some specific conditions of efficiency are not met, the forecast displaying the lowest MSPE will also display the lowest correlation with the target variable. Given that violations of efficiency are usual in the forecasting literature, this opposite behavior in terms of accuracy and correlation with the target variable may be a fairly common empirical finding that we label here as "the MSPE Paradox." We characterize "Paradox zones" in terms of differences in correlation with the target variable and conduct some simple simulations to show that these zones may be non-empty sets. Finally, we illustrate the relevance of the Paradox with two empirical applications.

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

**Keywords**: Mean Squared Prediction Error, Correlation, Forecasting, Time Series, Random Walk.

[1] Corresponding Author: Pablo Pincheira-Brown: Diagonal Las Torres 2640, Peñalolén, Santiago, Chile. Email: pablo.pincheira@uai.cl.

[♣] We are thankful to Universidad Adolfo Ibáñez for financial support through the Concurso de Investigación Individual 2020. We are also thankful to Juan Pablo Medina, Rodrigo Wagner and participants of the seminars of the Business School, Universidad Adolfo Ibáñez, for very constructive comments.

[2] The full version of the paper is available upon request: pablo.pincheira@uai.cl, nhardy@uft.cl

## 1. Introduction

"*How wonderful that we have met with a paradox. Now we have some hope of making progress*." Niels Bohr.

In this paper, we show that traditional comparisons of Mean Squared Prediction Error (MSPE) between two competing forecasts may be highly controversial. This is so because when some specific conditions of efficiency are not met, the forecast displaying the lowest MSPE will also display the lowest correlation with the target variable. Given that violations of efficiency are usual in the forecasting literature, this opposite behavior in terms of accuracy and correlation with the target variable may be a fairly common empirical finding that we label here as the MSPE paradox[3].

It is safe to say that MSPE is one of the most popular measures in the forecast evaluation literature, with a long tradition in both empirical and theoretical works. Just as an anecdotal illustration of its relevance, the acronym "MSPE" is mentioned 77 times in West's (2006) survey. The rationale for using MSPE as a loss function is as follows: MSPE is a statistical measure of accuracy, then a forecast displaying a low MSPE is an accurate forecast that, on average, will be close to the target variable. Some of the most iconic empirical contributions in economic forecasting (such as those of Meese and Rogoff (1983, 1988), Goyal and Welch (2008), Stock and Watson (2003), and Timmermann (2008)) rely partially or completely on MSPE comparisons. Due to its importance and tractability, it is not surprising that many theoretical works in this literature focus on MSPE as a leading case (e.g., Diebold and Mariano (1995), West (1996), Giacommini and White (2006)).

An alternative avenue to evaluate predictive ability could consider the association between the forecast and the target variable: the tighter the association is, the better the forecast is. Probably the simplest association measure between two random variables, X and Y, is the correlation between them. According to this intuition, a forecast more closely related to Y would be superior to another forecast not as closely related to Y. In other words, a forecast with a higher correlation with Y should be preferable to another forecast displaying a lower correlation.

Interestingly, in this paper, we show analytically and empirically that the forecast with the lowest MSPE does not necessarily display the highest correlation (what we call the MSPE Paradox). We show that both approaches are equivalent when the forecasts meet some conditions of efficiency (Mincer and Zarnowitz (1969)). Given that violations of efficiency

[3] For instance, Ince and Molodtsova (2017) study surveys of exchange rates expectations. They reject the null hypothesis of efficiency for nine developed economies (plus the euro area), and for most of the 23 emerging economies included in their analysis.

are usual in the forecasting literature (e.g., Ince and Molodtsova (2017)), this opposite behavior in terms of accuracy and correlation with the target variable may be a fairly common empirical finding.

We offer a characterization of “Paradox zones” in terms of the differences in correlation with the target variable. Moreover, we carry out simple simulations to show that these Paradox zones are, in general, not empty sets. As a matter of fact, our analysis shows that we could have an extreme case in which a totally uncorrelated forecast with the target variable could be superior in terms of MSPE to an alternative forecast displaying a positive correlation with the same target variable. Our empirical illustration supports this idea.

Finally, we show the relevance of the MSPE Paradox with two empirical applications in which some of the most accurate forecasts are, in fact, the worst in terms of correlations with the target variable. Both illustrations are related to the commodity-currencies literature. In the first exercise, we predict the returns of eleven commodities with the exchange rates of five commonly studied commodity-exporting economies. In the second exercise, we evaluate several exchange rates forecasts of the same five commodity-exporting economies. In this case, we compare the predictions of the FX4cast survey with some forecasts constructed with commodity returns and some usual benchmarks as well.

The rest of this paper is organized as follows. In section 2 we show with simple examples what we call the MSPE Paradox. We warn the reader that in subsection 2.1 we will be making very restrictive assumptions for the sake of clarity. Nevertheless, in subsection 2.2 we relax these assumptions to analyze the Paradox with more generality. In Section 3 we offer a characterization of the “Paradox zones”. In section 4 we illustrate the Paradox with simple simulations whereas in section 5 we present two empirical illustrations. Finally, section 6 concludes.

## 2. The MSPE Paradox

In this summary, we show simple examples of what we call "The MSPE Paradox." For a detailed discussion, see the full version of the paper (available upon request).[4]

### 2.1 Simple examples

In this section we illustrate with simple examples what we call "The MSPE Paradox." We use this name to label the fact that when comparing two competing forecasts for the same target variable, it might be the case that the forecast displaying the lowest MSPE will also display the lowest correlation with the target variable.

[4] The full version of the paper is available upon request: pablo.pincheira@uai.cl, nhardy@uft.cl

Let us consider $\{Y_t\}$ to be a mean zero target variable with variance equal 1. At time $t$, we have two competing forecasts $\{X_{t-1}\}$ and $\{Z_{t-1}\}$ for $\{Y_t\}$. It is important to notice that both $\{X_{t-1}\}$ and $\{Z_{t-1}\}$ are forecasts constructed with information previous to time $t$ and that they are taken as primitives (hence, we are not concerned here about parameter uncertainty). For clarity of exposition, we drop the sub-indexes $t$ in what follows. Let us assume that the vector $(Y, X, Z)'$ is weakly stationary (so here we assume the existence of second moments).

***Example 1:***

For example 1 we will also assume that both forecasts have the same non-negligible variance: Var(X)=Var(Z)>0, that X is a mean zero forecast and that $E(X^2) > 0$. Many of these assumptions are very restrictive, but they are useful to illustrate the Paradox.

Consider now the MSPE of both forecasts:

$$MSPE_X = E(Y - X)^2 \; ; \; MSPE_Z = E(Y - Z)^2$$

And let us also define the corresponding Mean Squared Forecasts (MSF) as follows:

$$MSF_X = E(X^2); \; MSF_Z = E(Z^2)$$

Suppose now that we are interested in a traditional comparison of MSPE, then:

$$\Delta MSPE \equiv MSPE_X - MSPE_Z = E(Y - X)^2 - E(Y - Z)^2$$

$$= (EX^2 - EZ^2) - 2(EYX - EYZ)$$

$$= (MSF_X - MSF_Z) - 2\{Cov(Y, X) - Cov(Y, Z)\}$$

$$= (MSF_X - MSF_Z) - 2\sqrt{Var(Y)}\left\{Corr(Y, X)\sqrt{Var(X)} - Corr(Y, Z)\sqrt{Var(Z)}\right\}$$

$$= (MSF_X - MSF_Z) - 2\sqrt{Var(X)}\{Corr(Y, X) - Corr(Y, Z)\}$$

$$= (MSF_X - MSF_Z) - 2\sqrt{MSF_X - (EX)^2}\{Corr(Y, X) - Corr(Y, Z)\}$$

$$= (MSF_X - MSF_Z) - 2\sqrt{MSF_X}\{Corr(Y, X) - Corr(Y, Z)\} \quad (1)$$

Eq.(1) illustrates an important result: the difference in MSPE depends not only on the correlation between the forecasts with the target variable, but also on the MSF that are not directly linked to properties of the target variable. The problem in this illustration relies on a "magnitude" effect associated to the term $(MSF_X - MSF_Z)$: A high MSF of a forecast could more than offset its high correlation with the target variable and therefore the forecast itself could be outperformed by another less informational forecast with a lower MSF. In other words, in this example, traditional MSPE comparisons give a natural advantage to "small forecasts", that is to say, forecasts with small MSF.

***Example 2:***

As a second example, let us consider a different econometric context, but similarly simplistic, in which Z is a zero-forecast. Consequently, $Var(Z) = Cov(Y,Z) = EZ^2 = 0$. Furthermore, let us also assume that $EX = 0$, Var(X) > 0 and that Var(Y) = 1. We will have then

$$\Delta MSPE \equiv MSPE_X - MSPE_Z = E(Y-X)^2 - E(Y-Z)^2$$

$$= (EX^2 - EZ^2) - 2(EYX - EYZ)$$

$$= (EX^2 - EZ^2) - 2\{Cov(Y,X) - Cov(Y,Z)\}$$

$$=(EX^2) - 2\{Cov(Y,X)\}$$

$$= (EX^2) - 2\sqrt{Var(Y)}\sqrt{Var(X)}\{Corr(Y,X)\}$$

$$= MSF_X - 2\sqrt{MSF_X}\{Corr(Y,X)\}$$

then if $\sqrt{MSF_X} > 2\,Corr(Y,X) > 0$ we will have that $MSPE_X > MSPE_Z$ despite that $Cov(Y,X) > Cov(Y,Z) = 0$. This is, of course, an extreme situation. The use of MSPE in this case, will suggest that the forecast with no association whatsoever with the target variable is preferable to another forecast with a tighter association. The problem in this example is that MSPE comparisons would fail to detect the usefulness of forecast X if its magnitude ($MSF_X$) overshadows its informational content.

**2.2 A general case**

In the full version of this paper (available upon request), we show a general decomposition of $\Delta MSPE$ (leaving behind the previous simplifying assumptions). We focus on two different scenarios: one in which Z is just a constant, and the other in which Z is a forecast with positive variance. We will refer to these conditions as C1 and C2:

C1) Z is just a constant *c*.

C2) Z is a forecast with positive variance.

The general decomposition for case C2 is

$$= MSF_X - MSF_Z - 2\sqrt{V(Y)}\left\{Corr(Y,X)\sqrt{MSF_x - (EX)^2} - Corr(Y,Z)\sqrt{MSF_z - (EZ)^2}\right\} - 2\{EY(EX - EZ)\} \quad (2.1)$$

While the general decomposition for case C1 is simply

$$\Delta MSPE = MSF_X - MSF_Z - 2Corr(Y,X)\sqrt{V(Y)V(X)} - 2EY(EX - EZ) \quad (2.2)$$

Notice that our decomposition relates to Clark and West (2006, 2007) in the following sense: In the context of out-of-sample comparisons of nested models, Clark and West (2006, 2007) notice that under the null of equal population MSPE, the sample MSPE of the nesting model is expected to be higher than that of the nested one. The intuition is that the nesting model introduces noise into its forecasts through the estimation of parameters that, under the null, are equal to zero. This effect inflates the sample MSPE of the model with additional parameters. Our decomposition resembles the findings by Clark and West, and to some extent, it is even more general. Both Clark and West and us similarly argue that a plain look at MSPE comparisons may be misleading in some circumstances, given that they can be affected by several distortions. In the case of Clark and West, those distortions arise from parameter estimation error. In our approach, these distortions arise at the population level by comparing apple and oranges: forecasts with very different magnitude effects or very different biases. In other words, even at the population level, we may observe that some forecasts have a natural advantage in terms of MSPE relative to others, despite of being far less informational relative to its competitors.

## 3. Some simple theoretical results

In the following, we will assume, without loss of generality, that $Corr(Y,X) \geq Corr(Y,Z)$ if Z has a positive variance. In case that Z has zero variance, we will assume without loss of generality that $Cov(Y,X) \geq Cov(Y,Z) = 0$. In this setup the Paradox will exist whenever $MSPE_X - MSPE_Z > 0$. We will show analytically that the Paradox may exist under some conditions if the difference between $Corr(Y,X)$ and $Corr(Y,Z)$ is relatively small. As we are considering the two leading cases of positive and zero variance for Z we will denote by $\Omega_1$ the variance-covariance matrix of the (Y,X,Z)' vector and by $\Omega_2$ the variance-covariance matrix of the (Y,X)' vector:

**<u>Proposition 1:</u>** Let Z be a constant-forecast (say, $Z = c\ \forall t$). Let us also assume that

$$\frac{MSF_X - c^2}{2\sqrt{V(Y)V(X)}} - \frac{EY(EX-c)}{\sqrt{V(Y)V(X)}} > 0$$

Then we will find the Paradox iif $Corr(Y,X) \in [0; \frac{MSF_X - c^2}{2\sqrt{V(Y)V(X)}} - \frac{EY(EX-c)}{\sqrt{V(Y)V(X)}})$.

**Corollary 1:** If in Proposition 1 we set $Z = c = 0$, we will find the Paradox iif $Corr(Y,X) \in [0; \frac{MSF_X}{2\sqrt{V(Y)V(X)}} - \frac{EYEX}{\sqrt{V(Y)V(X)}})$. Moreover, in the particular case in which $EX = 0$, we will find the Paradox iif $Corr(Y,X) \in [0; \frac{\sqrt{V(X)}}{2\sqrt{V(Y)}})$.

**Proof of Proposition 1.**

The proofs of Proposition 1 and Corollary 1 are available in the full version of this paper (upon request). ■

Proposition 2 next shows an equivalent result for the case in which Z is a forecast with positive variance.

**Proposition 2:** Let Z be a forecast with positive variance. Let $\Delta = Corr(Y,X) - Corr(Y,Z)$ and suppose that

$$\frac{MSF_X - MSF_Z}{2\sqrt{V(Y)V(X)}} - \frac{Corr(Y,Z)\left(\sqrt{V(X)} - \sqrt{V(Z)}\right)}{\sqrt{V(X)}} - \frac{EY(EX - EZ)}{\sqrt{V(Y)V(X)}} > 0$$

Then we will find the Paradox iif

$$\Delta \in [0; \frac{MSF_X - MSF_Z}{2\sqrt{V(Y)V(X)}} - \frac{Corr(Y,Z)\left(\sqrt{V(X)} - \sqrt{V(Z)}\right)}{\sqrt{V(X)}} - \frac{EY(EX - EZ)}{\sqrt{V(Y)V(X)}}).$$

**Proof of Proposition 2.**

The proof of Proposition 2 is available in the full version of this paper (upon request). ■

Next we will see that for the Paradox to exist we require inneficient forecasts. We need some notation first: Let $\mathbf{u_x}$ and $\mathbf{u_z}$ be the forecast errors of X and Z, respectively. In other words

$$\mathrm{u_x} \equiv Y - X$$

$$\mathrm{u_Z} \equiv Y - Z$$

Let us recall that X and Z are efficient forecasts à la Mincer and Zarnowitz (1969) as long as

$$Cov(X, u_x) = Cov(Z, u_z) = 0$$

$$E(u_x) = E(u_z) = 0$$

**Proposition 3:** If X and Z are both efficient *à la* Mincer and Zarnowitz, then the Paradox is impossible.

**Proof of Proposition 3.**

The proof of Proposition 3 is available in the full version of this paper (upon request). ■

## 4. Simulations

To illustrate the Paradox and to show that the "Paradox zone" may be a non-empty set under inefficiencies, we carry out a set of simple simulations. In each simulation, we show that the Paradox zone coincides with the intervals derived in Section 3. In the full version of the paper (upon request), we show three different simulations: i) Z as a zero-forecast with arbitrary parameters, ii) Z is a more general case with arbitrary parameters, and iii) Z is a zero-forecast in the context of a data generating process calibrated with exchange rates forecasts. In this summary, we only show results for simulation i).

### 4.1 Simulation with a "zero-forecast."

Let us suppose we wish to compare two competing forecasts, X and Z, where Var(X) > 0 and Z is a "zero-forecast." According to corollary 1 in Section 3, the Paradox zone is defined by

$$Corr(X,Y) \in [0; \frac{MSFx}{2\sqrt{V(Y)V(X)}} - \frac{E(Y)E(X)}{\sqrt{V(Y)V(X)}})$$

To show that $[0; \frac{MSFx}{2\sqrt{V(Y)V(X)}} - \frac{E(Y)E(X)}{\sqrt{V(Y)V(X)}})$ may be a non-empty region, we consider the following simulation: We start by setting $EY = 0.1$, $EX = 1$, $E(Y^2) = 2$ and $E(X^2) = 2$, then from expression (4) we have

$$\Delta MSPE = \text{MSPE}_\text{x} - MSPE_z = 1.8 - 2.8213 * Corr(X,Y)$$

Keeping $EY, EX, E(Y^2)$ and $E(X^2)$ constant, our decomposition is just a linear function between $\Delta MSPE$ and Corr(X,Y), with a slope of -2.8213 and an intercept of 1.8. In order to analyze this linear function without changing the slope nor the intercept, we generate different values of $Corr(X,Y)$ just by changing EYX but keeping in mind that the covariance matrix $\Omega_2$ must be positive definite:

$$\Omega_2 = \begin{pmatrix} E(X^2) - (EX)^2 & EYX - EYEX \\ EYX - EYEX & E(Y^2) - (EY)^2 \end{pmatrix} = \begin{pmatrix} 1 & EYX - 0.1 \\ EYX - 0.1 & 1.99 \end{pmatrix}$$

We parameterize $\text{EYX} = 0.1 + \delta$, where $\delta$ is a sequence of small positive incremental changes of 0.001. Notice that the slope and the intercept of our linear function remain unaltered in this simulation. In this case, the Paradox zone is given by $Corr(X,Y) \in \left[0; \frac{E(X^2)}{2\sqrt{V(Y)V(X)}} - \frac{E(Y)E(X)}{\sqrt{V(Y)V(X)}}\right[ = [0; 0.638)$. In other words, despite that forecast Z has no covariance with Y, it outperforms the forecast X in terms of MSPE whenever $Corr(X,Y) \in [0; 0.638)$.

**Figure 1: Illustration of the Paradox zone when Z is a zero-forecast.**

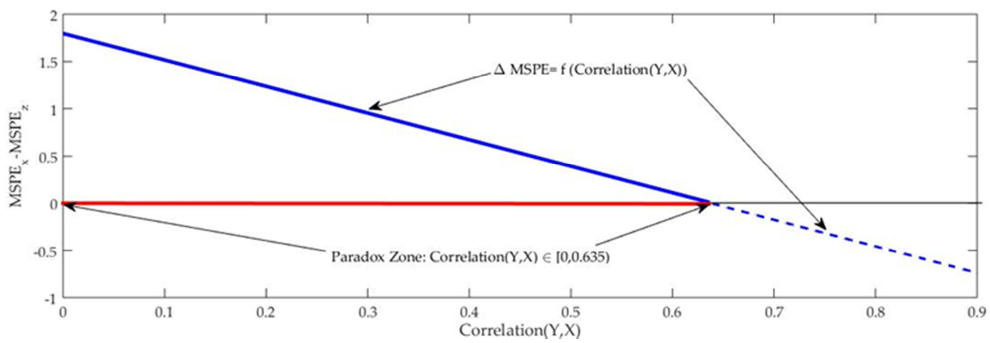


Source: Author's elaboration

### 4.2 Simulation with a general forecast Z

Available in the full version of the paper (upon request).

### 4.3 Simulation calibrated with exchange rates.

Available in the full version of the paper (upon request).

## 5. Empirical illustrations of the MSPE Paradox

In the full version of the paper (available upon request), we illustrate the Paradox with two empirical applications using commodities and commodity currencies. In both cases we will assume for simplicity that population moments are well approximated by their sample counterparts. In this summary, we show exclusively one empirical illustration forecasting commodity currencies.

### 5.1 The Paradox in commodity forecasts

Our first empirical illustration is mainly inspired by the commodity-currencies literature. Chen, Rogoff and Rossi (2010, 2011) seminal papers report strong predictive ability from the exchange rates of some exporting countries such as Australia, Canada, Chile, New Zealand and South Africa to some country-specific commodity indices. Additionally, Pincheira and Hardy (2019a, 2019b) find strong predictive ability from the Chilean peso to base-metal prices and the London Metal Exchange Index (LMEX).

In this context, we construct and compare different forecasts for 11 series of commodities (aluminum, copper, lead, nickel, zinc, tin, LMEX, gold, silver, S&P GSCI, and platinum) using the exchange rates of Australia, Canada, Chile, New Zealand and South Africa (relative to the U.S dollar). The econometric specifications for our forecasts closely follow Pincheira and Hardy (2019a, 2019b):

$$\Delta CP_t = \beta \Delta E\mathrm{R}_{\mathrm{t-1}} + \varepsilon_{\mathrm{t}} \;\; (M1)$$

Where $\Delta CP_t$ stands for the log-difference of a commodity price, $\Delta ER$ is the log-difference of a generic exchange rate, $\beta$ is a regressor coefficient and $\varepsilon_{\mathrm{t}}$ is the error term. Note that we are only evaluating one-step-ahead forecasts. The database is collected from Thomson Reuters Datastream, considering monthly closing prices on commodity prices and exchange rates (relative to the U.S dollar). In this analysis, we consider exclusively a period in which all the economies pursue a pure flotation exchange rate regime; hence our database goes from October 1999 through May 2019 (a total of T=236 observations for each series).

In addition to the five forecasts generated by each exchange rate using (M1), we also consider the forecasts of a Driftless Random Walk (a zero-forecast, DRW), a Random Walk (a forecast with the historical mean, RW) and an AR(1). Finally, the parameter $\beta$ in (M1), and the parameters of the AR(1) and the RW are estimated by OLS and updated with rolling windows of R=48 monthly observations. Notice that all our forecasts are evaluated out-of-sample, with a total of P=T-R=188 predictions.

**Table 1: Evaluation of commodity forecast with correlations and RMSPE**

| *(1)* | *(2)* | *(3)* | *(4)* | *(5)* | *(6)* | *(7)* | *(8)* | *(9)* |
|---|---|---|---|---|---|---|---|---|
| | Australia | Canada | Chile | NewZealand | SouthAfrica | AR(1) | RW | DRW |
| | | | | S&P GSCI | | | | |
| Correlation | 0.070 | -0.023 | 0.147 | -0.118 | 0.031 | 0.151 | -0.065 | - |
| RMSPE | 0.067 | 0.067 | 0.066 | 0.069 | 0.067 | 0.066 | 0.067 | 0.066 |
| | | | | LMEX | | | | |
| Correlation | 0.049 | -0.025 | 0.146 | -0.178 | -0.046 | 0.161 | 0.009 | - |
| RMSPE | 0.068 | 0.067 | 0.067 | 0.069 | 0.068 | 0.067 | 0.067 | 0.066 |
| | | | | Aluminum | | | | |
| Correlation | 0.029 | -0.068 | 0.139 | 0.007 | -0.087 | 0.137 | -0.043 | - |
| RMSPE | 0.063 | 0.063 | 0.062 | 0.063 | 0.063 | 0.063 | 0.063 | 0.062 |

Notes: RMSPE stands for Root MSPE. Source: Author's elaboration

Table 1 reports our results when predicting GSCI, LMEX and aluminum[5]. First, 12 out of 21 non zero forecasts have a positive correlation with the target variable, suggesting some useful information in those forecasts. Notably, the DRW is the forecast with the lowest RMSPE in the three commodities, despite having zero covariance with the target variable.

Second, note that the forecasts for each commodity show very similar RMSPE, but very different correlations. For instance, the RMSPE for the LMEX goes between 0.066 and 0.069, but the correlations vary between -0.178 and 0.161. In other words, relative to the maximums we find changes of around 4% in RMSPE and changes of around 210% for correlations.

Third, there are some cases of paradoxes worth to be mentioned. For instance, for the LMEX, the forecast of the AR(1) has a particularly high correlation of 0.161; nevertheless, the DRW has a lower RMSPE. Moreover, the forecast constructed with the Australian dollar has a correlation of 0.049, but notably, it has greater RMSPE than the forecast constructed with the Canadian dollar, even though the latter has a negative correlation of -0.025.

To illustrate the Paradox, Figures 4 and 5 display the differences in MSPE and correlations between two forecasts using rolling windows of 48 observations. Figure 4 compares two different forecasts for aluminum: one constructed with the Australian Dollar, the other with the South African Rand. Figure 5 reports our results when forecasting the LMEX with the Australian and the Canadian Dollar. In both figures, whenever the differences in MSPE and correlations have the same sign, we have the MSPE Paradox. Notably, both figures suggest that the Paradox may appear quite often.

[5] Results for the rest of commodities are presented in the Appendix of the full version of the paper (available upon request).

**Figure 4: Differences in MSPE and Correlations using rolling windows. Forecasting aluminum with the Australian and South African exchange rates.**

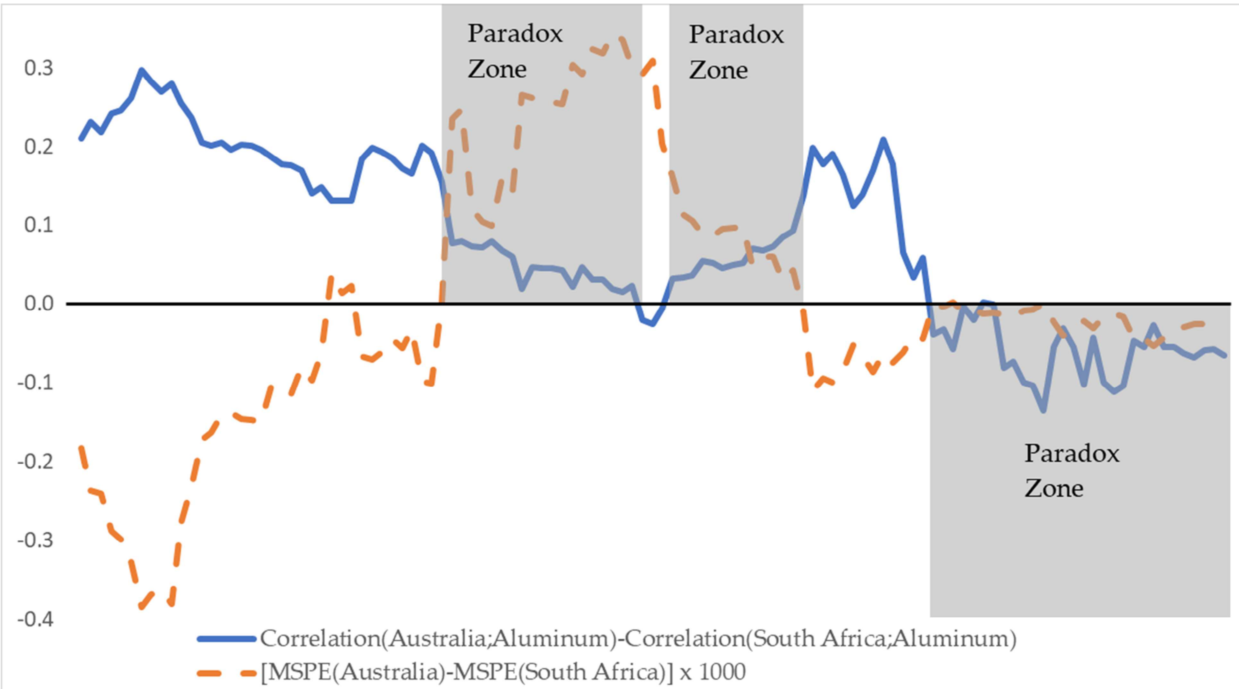


Notes: Figure 4 displays the differences in MSPE and correlations between two competing forecasts, using rolling windows of 48 observations. In this illustration the target variable is aluminum one-month returns. We compare a forecast using the Australian Dollar with another using the South African Rand. Whenever both series have the same sign, we have the MSPE Paradox. The differences in MSPE have been scaled so the left axis represents both differences in MSPE and differences in correlations. Source: Author's elaboration

**Figure 5: Differences in MSPE and Correlations using rolling windows. Forecasting LMEX with the Australian and Canadian exchange rates.**

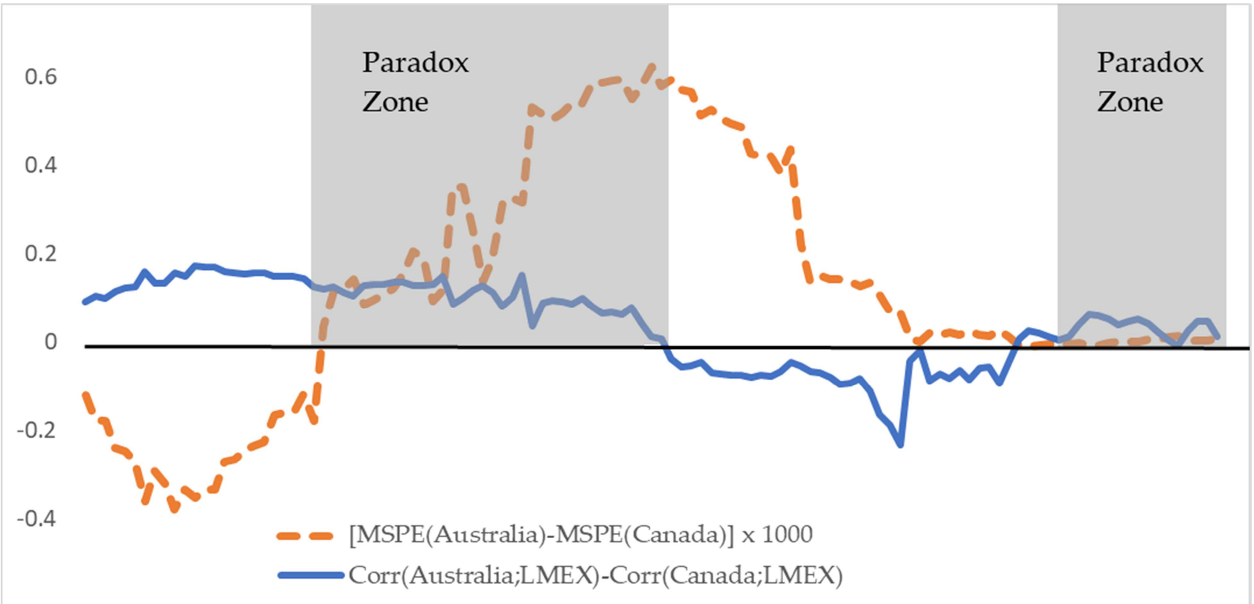


Notes: Figure 5 displays the differences in MSPE and correlations between two competing forecasts using rolling windows of 48 observations. In this illustration the target variable is LMEX one-month returns. We compare a forecast using the Australian Dollar with another using the Canadian Dollar. Whenever both series have the same sign, we have the MSPE Paradox. The differences in MSPE have been scaled so the left axis represents both differences in MSPE and differences in correlations. Source: Author's elaboration

### 5.2 The Paradox in exchange rates forecasts

Available in the full version of the paper (upon request).

## 6. Concluding remarks

In this paper we show that traditional comparisons of Mean Squared Prediction Error (MSPE) between two competing forecasts may be highly controversial. This is so because when some specific conditions of efficiency are not met, the forecast displaying the lowest MSPE will also display the lowest correlation with the target variable. Given that violations of efficiency are usual in the forecasting literature, this opposite behavior in terms of accuracy and correlation with the target variable may be a fairly common empirical finding that we label here as "the MSPE Paradox."

We characterize "Paradox zones" in terms of differences in correlation with the target variable and conduct some simple simulations to show that these zones may be non-empty sets. Moreover, our analysis shows that we could have an extreme case in which a forecast and a target variable are independent random variables, which speaks of a useless forecast, yet, in terms of MSPE, this useless forecast might outperform a useful forecast displaying a positive correlation with the target variable.

Finally, we illustrate the relevance of the MSPE Paradox with two empirical applications in which some of the most accurate forecasts in terms of MSPE are, in fact, some of the worst in terms of correlations with the target variable.

Our paper emphasizes the need to look beyond MSPE when evaluating two or more competing forecasts, as a blind search for the minimum out-of-sample MSPE forecast may lead to an incorrect evaluation of the information contained within those predictions. In light of these results, an interesting avenue for future research is the elaboration of a simple asymptotically normal test to evaluate two competing forecasts according to their correlations with the target variable.

## 7. References


1. Chen, Y.-C., Rogoff, K. S., and Rossi, B. (2010). Can Exchange Rates Forecast Commodity Prices ? *Quarterly Journal of Economics*, 125(August), 1145–1194.
2. Chen, Y.-C., Rogoff, K. S., and Rossi, B. (2011). Predicting Agri-Commodity Prices: An Asset Pricing Approach, World Uncertainty and the Volatility of Commodity Markets, ed. B. Munier, IOS.
3. Clark, T. E., and West, K. D. (2006). Using out-of-sample mean squared prediction errors to test the martingale difference hypothesis. *Journal of econometrics*, *135*(1-2), 155-186.
4. Clark, T. E., and West, K. D. (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of econometrics*, *138*(1), 291-311.
5. Diebold, F. X., and Mariano, Roberto, S. (1995). Comparing Predictive Accuracy. *Journal of Business and Economic Statics*, 13(3), 253–263.
6. Ferraro, D., Rogoff, K., and Rossi, B. (2015). Can oil prices forecast exchange rates? An empirical analysis of the relationship between commodity prices and exchange rates. *Journal of International Money and Finance*, 54, 116-141.
7. Giacomini, R., and White, H. (2006). Tests of conditional predictive ability. *Econometrica*, *74*(6), 1545-1578.

8. Ince, O., and Molodtsova, T. (2017). Rationality and forecasting accuracy of exchange rate expectations: Evidence from survey-based forecasts. *Journal of International Financial Markets, Institutions and Money*, *47*, 131-151.
9. Meese, R, and Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics*, 14, 3–24.
10. Meese, R., and Rogoff, K. (1988). Was it real? The exchange rate-interest differential relation over the modern floating-rate period. *the Journal of Finance*, *43*(4), 933-948.
11. Mincer, J. A., and Zarnowitz, V. (1969). The evaluation of economic forecasts. In *Economic forecasts and expectations: Analysis of forecasting behavior and performance* (pp. 3-46). NBER.
12. Pincheira, P., and Hardy, N. (2019a). Forecasting base metal prices with the Chilean exchange rate. *Resources Policy*, 62(February), 256–281.
13. Pincheira, P., and Hardy, N. (2019b). Forecasting Aluminum Prices with Commodity Currencies Forecasting Aluminum Prices with Commodity Currencies. Available at SSRN 3511564.
14. Stock, J. H., and Watson, M. (2003). Forecasting output and inflation: The role of asset prices. *Journal of Economic Literature*, 41(3), 788-829.
15. Timmermann, A. (2008). Elusive return predictability. *International Journal of Forecasting*, *24*(1), 1-18.
16. Welch, I., and Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies*, 21(4), 1455-1508.
17. West, K. D. (1996). Asymptotic Inference about Predictive Ability. *Econometrica*, 64(5), 1067.
18. West, K. D. (2006). Chapter 3 Forecast Evaluation. Handbook of Economic Forecasting, 1(05), 99–134.