

CAMP Working Paper Series  
No 6/2020

# Proper scoring rules for evaluating asymmetry in density forecasting

Matteo Iacopini, Francesco Ravazzolo and Luca Rossini



© Authors 2020 This paper can be downloaded without charge from the CAMP website [bi.no/camp](http://bi.no/camp)



# Proper scoring rules for evaluating asymmetry in density forecasting\*

Matteo Iacopini<sup>†</sup>

Francesco Ravazzolo<sup>‡</sup>

Luca Rossini<sup>§</sup>

September 1, 2020

## Abstract

This paper proposes a novel asymmetric continuous probabilistic score (ACPS) for evaluating and comparing density forecasts. It extends the proposed score and defines a weighted version, which emphasizes regions of interest, such as the tails or the center of a variable's range. A test is also introduced to statistically compare the predictive ability of different forecasts. The ACPS is of general use in any situation where the decision maker has asymmetric preferences in the evaluation of the forecasts. In an artificial experiment, the implications of varying the level of asymmetry in the ACPS are illustrated. Then, the proposed score and test are applied to assess and compare density forecasts of macroeconomic relevant datasets (US employment growth) and of commodity prices (oil and electricity prices) with particular focus on the recent COVID-19 crisis period.

**Keyword:** asymmetric continuous probabilistic score; asymmetric loss; proper score; density forecast; predictive distribution; weighted score; probabilistic forecast.

## 1 Introduction

Macroeconomic forecasting has always been of pivotal importance for central bankers, policymakers and researchers. Nowadays, the vast majority of the research in macroeconomics and finance mainly focuses on the development and implementation of forecasting techniques that minimize the expected squared forecast error (Gneiting (2011)). This approach is grounded on the implicit assumption of using a symmetric loss function in evaluating the accuracy of a forecast.

Despite being common practice, the use of symmetric loss functions in forecasting is unrealistic especially in policy institutions, where the policymakers could have a specific

---

\*The authors gratefully acknowledge Todd Clark, Michael McCracken, Massimiliano Marcellino, Barbara Rossi, Jonas Brehmer for their useful feedback. This paper is part of the research activities at the Centre for Applied Macroeconomics and Commodity Prices (CAMP) at the BI Norwegian Business School. This research used the SCSCF multiprocessor cluster system at Ca' Foscari University of Venice. Luca Rossini acknowledges financial support from the EU Horizon 2020 programme under the Marie Skłodowska-Curie scheme (grant agreement no. 796902).

<sup>†</sup>Ca' Foscari University of Venice, Italy. [matteo.iacopini@unive.it](mailto:matteo.iacopini@unive.it)

<sup>‡</sup>Free University of Bozen-Bolzano, Italy and CAMP, BI Norwegian Business School, Norway. [francesco.ravazzolo@unibz.it](mailto:francesco.ravazzolo@unibz.it)

<sup>§</sup>Queen Mary University of London, United Kingdom and Vrije Universiteit Amsterdam, The Netherlands. [l.rossini@qmul.ac.uk](mailto:l.rossini@qmul.ac.uk)

aversion to positive or negative deviations of a forecast from the target. Consider a policymaker who is interested in forecasting employment. Suppose that, if the predicted employment rate drops below a given threshold, she will be forced to adopt new expansionary economic policy. It is highly likely that the policymaker is more averse to forecasts that give too high probability mass to the right part of the distribution of the employment rate (positive growth of employment), while she may be more relaxed with respect to forecasts that give too high probability mass to the left part of the distribution (negative or low growth of employment). Other examples relate to energy markets that have recently experienced negative prices. WTI oil prices collapsed to -37.63 US dollar for barrel in April 2020; German electricity prices have measured several negative prices with the introduction of renewable energy resources (RES). Producers would be more sensitive to prices below a threshold, up to zero if the marginal cost of production is zero, as it is the case of RES, than higher prices. These examples call for the design of a more general class of loss functions and scoring rules that account for asymmetry, in order to guide the process of making and assessing forecasts. To the best of our knowledge, a measure that properly incorporates asymmetry in density forecasting evaluation does not exist in the literature.

The main goal of this paper is the proposal of novel and practical forecasting evaluation tools that can fill in this gap and answer the increasing demand from policymakers and central bankers. We plan to achieve this result by introducing an innovative asymmetric scoring rule that is able to measure and evaluate heterogeneous aversion to different deviations of a density forecast from the target. We derive some properties of the new scoring rule and in particular demonstrate that it is a proper scoring rule. Moreover, we provide threshold- and quantile-weighted versions that allow to emphasize the performance of the forecast in regions of interest to the policymaker.

Within the literature on point forecasting, [Christoffersen and Diebold \(1996, 1997\)](#) proposed some asymmetric loss functions. In the former paper, they studied the optimal prediction problem under general loss structures and characterized the optimal predictor under an asymmetric loss function, focusing on the *LinEx* and the *LinLin* asymmetric functions. In the latter paper, the authors provided an illustration of an asymmetric loss in the context of GARCH processes.

More recently, scholars have begun to empirically investigate the degree of loss function asymmetry of central banks and other international institutions. Among others, [Elliott et al. \(2005\)](#) and [Patton and Timmermann \(2007\)](#) proposed formal methods to infer the degree of asymmetry of the loss function and to test the rationality of forecasts. Within this stream of literature, [Artis and Marcellino \(2001\)](#) found that IMF and OECD forecasts of the deficit of G7 countries are biased towards over or under-prediction relative to mean square error (MSE) forecasts. Regarding European institutions forecasts, [Christodoulakis and Mamatzakis \(2008, 2009\)](#) found evidence of asymmetric loss. In another study, [Dovern and Janssen \(2017\)](#) documented that the GDP growth forecasts made by professional forecasters tend to exhibit systematic errors, and tend to overestimate GDP growth. Moreover, [Boero et al. \(2008\)](#)

interpreted the tendency to over-predict GDP growth as a signal that policymakers exhibit greater fear of under-prediction than over-prediction, thus suggesting that their judgements are based on an asymmetric loss. More recently, [Tsuchiya \(2016\)](#) examined the asymmetry of the loss functions of the Japanese government, the IMF and private forecasters for Japanese growth and inflation forecasts.

In the framework of forecast combination, [Elliott and Timmermann \(2004\)](#) showed that the optimal combination weights significantly differ under asymmetric loss functions and skewed error distributions as compared to those obtained with mean squared error loss. Finally, [Demetrescu and Hoke \(2019\)](#) studied factor-augmented forecasting under asymmetric point loss function.

An alternative and more universal approach to forecasting is the provision of a predictive density, known as probabilistic or density forecasting (see [Elliott and Timmermann \(2016a, ch.8\)](#)). Two key aspects of density forecasts are the statistical compatibility between the forecasts and the realized observations (calibration) and the concentration of predictive distributions (sharpness). The aim of probabilistic forecasts is to maximize their sharpness, subject to calibration ([Gneiting and Ranjan \(2013\)](#)). Density forecasting is more complex than point forecasting since the estimation problem requires to construct the whole predictive distribution, rather than a specific function thereof (e.g., mean or quantile). Several reasons have been suggested for preferring density over point forecasts (e.g., [Elliott and Timmermann \(2016b\)](#)). First, point forecasting is often associated to the mean of a distribution and it is optimal for highly restricted loss functions, such as quadratic loss function, but inadequate for any prospective user having a different loss. Moreover, the value of a point forecast can be increased by supplementing it with some measures of uncertainty and complete probability distributions over the outcomes provide useful information for making economic decisions; see, for example, [Anscombe \(1968\)](#) and [Zarnowitz \(1969\)](#) for early works and the discussions in [Granger and Pesaran \(2000\)](#), [Timmermann \(2006\)](#) and [Gneiting \(2011\)](#). [Carriero et al. \(2020\)](#) extends the application to tail risk nowcasts of economic activity. Finally, in recursive forecasting with nonlinear models the full predictive density matters since the nonlinear effects typically depend not only on the conditional mean, but also on where future values occur in the set of possible outcomes.

A natural way to evaluate and compare competing density forecasts is the use of proper scoring rules, which assess calibration and sharpness simultaneously and encourage honest and careful forecasting. Despite the wide literature on the class of proper scoring rules for probabilistic forecasts of categorical and binary variables (e.g., see [Savage \(1971\)](#), [Schervish \(1989\)](#)) the advances for continuous variables are more limited. Motivated by these facts, we aim at designing a novel asymmetric proper scoring rule to be used for evaluating density forecasts of continuous variables, which is the typical case in macroeconomics and finance exercises (e.g., predicting variables such as unemployment, inflation, log-returns, GDP growth, and realized volatility).

[Gneiting and Raftery \(2007\)](#) proposed the continuous rank probability score (CRPS) as

a proper scoring rule for probabilistic forecasts of continuous variables, and more recently, [Gneiting and Ranjan \(2011\)](#) extended the CRPS by introducing a threshold- and a quantile-weighted version (tCRPS and qCRPS, respectively). These scoring rules give more emphasis to the performance of the density forecast in a selected *region of the domain*,  $B$ , by assigning more weight to the deviations from the observations made in  $B$ . The major drawback of both the CRPS and its weighted versions is the symmetry of the underlying reward scheme, meaning that they assign equal reward to positive and negative deviations of a probabilistic forecast from the target. This comes from the fact that the CRPS is built on the Brier score and inherits some of its properties, such as properness and symmetry. Similarly, since both the weighted versions of the CRPS essentially consist in re-weighting the CRPS over the domain of the variable of interest, they inherit the symmetry of the latter.

[Winkler \(1994\)](#) did a first effort towards asymmetric scoring rules and proposed a general method for constructing asymmetric proper scoring rules starting from symmetric ones. However, this approach is limited to forecasting binary variables, and continuous variables were not investigated.

We address this issue and contribute to the literature on proper scoring rules for evaluating density forecasts by proposing a novel asymmetric proper scoring rule which assigns different penalties to positive and negative deviations from the true density. The main contribution of this paper is twofold. First, we define a new proper scoring rule which assigns an asymmetric penalty to deviations from the target density. Moreover, we provide a threshold- and quantile-weighted version of it and develop an adaptation of the Diebold-Mariano test to statistically compare the predictive ability of different forecasts. Then, we compare the performance of the scores with the CRPS and its weighted versions. Second, we use the proposed score to evaluate density forecasts in three relevant applications in macroeconomics (US employment growth) and commodity prices (oil and electricity prices) with data updated to the COVID-19 crisis period. Variables have experienced large volatilities, with sizeable spikes and negative energy prices. As we discussed above, players might be more sensitive to some specific parts of the distribution of these series and we shed light on how to evaluate this asymmetry.

The key result of this paper is the provision of a tool able to account for the decision maker’s preferences in the evaluation of density forecasts, both in terms of domain- and error-weighting schemes. Domain-weighting gives heterogeneous emphasis to the performance on different regions, while the error-weighting asymmetrically rewards negative and positive deviations from the target value. The proposed weighted asymmetric scoring rule combines the two schemes and allows to evaluate the performance of the forecasting density from both perspectives.

The rest of the paper is organized as follows. [Section 2](#) presents a novel asymmetric scoring rule for density forecasts, its extension to threshold- and quantile-weighted versions and a test to compare the predictive accuracy of different forecasts. Then [Section 3](#) discusses its main properties. It also illustrates a comparison with the (weighted) CRPS in simulated experiments. Finally, [Section 4](#) provides different applications on forecasting US macroeconomic variables

(employment rate) and commodity prices (oil and electricity prices). The article closes with a discussion in Section 5.

The MATLAB code for implementing the proposed scoring rules is available at:

<https://github.com/matteoiacopini/acps>

## 2 Asymmetric Proper Scoring rules for Density forecasting

The evaluation and comparison of probabilistic forecasts typically relies on proper scoring rules. Informally, a scoring rule is a measure that summarises the goodness of a probabilistic forecast by combining the predictive distribution and the value that actually materializes. One can think of it as a measure of distance between the probabilistic forecast and the actual value. We consider positively oriented scoring rules, therefore if probabilistic forecast  $P_1$  obtains a higher score than  $P_2$ , this means that  $P_1$  yields a more accurate forecast than  $P_2$ . Therefore, the score can be interpreted as a reward to be maximized.

In more formal terms, following the notation of [Gneiting and Raftery \(2007\)](#), consider the problem of making probabilistic forecasts on a general sample space  $\Omega$ . Let  $\mathcal{A}$  be a  $\sigma$ -algebra of subsets of  $\Omega$ , and let  $\mathcal{P}$  be a convex class of probability measures on  $(\Omega, \mathcal{A})$ . A *probabilistic forecast* is any probability measure  $P \in \mathcal{P}$ , such that  $P : \Omega \rightarrow \bar{\mathbb{R}}$ , where  $\bar{\mathbb{R}} = [-\infty, +\infty]$  denotes the extended real line, is said to be  $\mathcal{P}$ -quasi-integrable if it is measurable with respect to  $\mathcal{A}$  and is quasi-integrable with respect to all  $P \in \mathcal{P}$  (see [Bauer \(2011\)](#)). A *scoring rule* is any extended real-valued function  $S : \mathcal{P} \times \Omega \rightarrow \bar{\mathbb{R}}$  such that  $S(P, \cdot)$  is  $\mathcal{P}$ -quasi-integrable for all  $P \in \mathcal{P}$ . In practice, if  $P$  is the forecast density and the event  $\omega$  materializes, then the forecaster's reward is  $S(P, \omega)$ .

In order to be effectively used in scientific forecasts evaluation, scoring rules have to be proper, meaning that they have to reward accurate forecasts. Suppose the true density of the observations is  $Q$  and denote the expected value of  $S(P, \omega)$  under  $Q(\omega)$  with

$$S(P, Q) = \mathbb{E}_Q[S(P, \omega)] = \int_{\Omega} S(P, \omega) Q(d\omega),$$

then the scoring rule  $S$  is *strictly proper* if  $S(Q, Q) \geq S(P, Q)$ . The equality holds if and only if  $P = Q$ , thus implying that the forecaster has higher reward if she predicts  $P = Q$ . If instead  $S(Q, Q) \geq S(P, Q)$  for all  $P$  and  $Q$ , then the scoring rule is said to be *proper*.

The vast majority of the proper scoring rules proposed in the literature are symmetric, that is, they reward in the same way positive and negative deviations from the target. For example, suppose a forecast  $P_1$  assigns too high probability mass to the right part of the domain (as compared to the true density) and a forecast  $P_2$  assigns too high probability mass to the left part, by the same amount. If these forecasts are evaluated under a symmetric scoring rule, then they receive the same score.

A symmetric loss is unsatisfactory for many real world situations where the decision maker has a preference or aversion towards a particular kind of error. We aim at filling in this gap



by defining a new asymmetric proper scoring rule for continuous variables, which is suited for evaluation and comparison of density forecasts and penalises more either side of the deviation from the target.

**Definition 1** (Asymmetric Continuous Probability Score). *Let  $c \in (0, 1)$  represent the level of asymmetry, such that  $c = 0.5$  implies a symmetric loss, while  $c < 0.5$  penalises more the left tail, and  $c > 0.5$  the right tail. Let  $P$  be the probabilistic forecast and  $y$  the realized (ex-post) value. We define the asymmetric continuous probability score (ACPS) as*

$$\begin{aligned} ACPS(P, y; c) = & \int_{-\infty}^y (c^2 - P(u)^2) \left[ \frac{1}{(1-c)^2} \mathbb{I}(P(u) > c) + \frac{1}{c^2} \mathbb{I}(P(u) \leq c) \right] du \\ & + \int_y^{+\infty} ((1-c)^2 - (1-P(u))^2) \left[ \frac{1}{(1-c)^2} \mathbb{I}(P(u) > c) + \frac{1}{c^2} \mathbb{I}(P(u) \leq c) \right] du. \end{aligned} \quad (1)$$

The following result shows the properness of our new score for every level of asymmetry.

**Theorem 1** (Properness). *The asymmetric scoring rule ACPS defined in eq. (1) is strictly proper for any  $c \in (0, 1)$ .*

*Proof.* The strict properness derives from the fact that ACPS can be obtained from the quadratic score for binary outcomes, which is strictly proper, via two transformations that preserve properness, see [Winkler \(1994\)](#) and [Matheson and Winkler \(1976\)](#). Specifically, let  $p \in (0, 1)$  be a probabilistic forecast of success in a binary experiment and let  $S$  be the quadratic rule, that is

$$S(p) = \begin{cases} S_1(p) = 1 - (1-p)^2, & \text{if success,} \\ S_2(p) = 1 - p^2, & \text{if failure.} \end{cases}$$

Notice that  $S(p)$  is a strictly proper and symmetric scoring rule. Following [Winkler \(1994\)](#), one can obtain a strictly proper asymmetric scoring rule for binary outcomes via the transformation

$$S_c^A(p) = \begin{cases} \frac{S_1(p) - S_1(c)}{T(c)}, & \text{if success,} \\ \frac{S_2(p) - S_2(c)}{T(c)}, & \text{if failure,} \end{cases} \quad T(c) = \begin{cases} S_1(1) - S_1(c), & \text{if } p > c, \\ S_2(0) - S_2(c), & \text{if } p \leq c, \end{cases}$$

where  $c \in (0, 1)$  denotes the level of asymmetry. Then, following [Matheson and Winkler \(1976\)](#), to obtain an asymmetric scoring rule for continuous variables, we assume that the subject assigns a probability distribution function  $P(x)$  to a continuous variable of interest. Fix an arbitrary real number  $u$  to divide the real line into two intervals,  $I_1 = \mathbb{I}(-\infty, u]$  and  $I_2 = \mathbb{I}(u, \infty)$ , and define a success the event that  $y$  falls in  $I_1$ . Since  $P(u) \in (0, 1)$  for any  $u \in \mathbb{R}$ , we can evaluate the binary scoring rule  $S_c^A$  at  $p = P(u)$ , thus obtaining a different value  $S_c^A(P(u))$  for each  $u$ . Finally, the dependence of the scoring rule on the arbitrary value of  $u$  is removed by integrating over all  $u$ , which yields eq. (1).

Notice that one can obtain a different (strictly) proper asymmetric scoring rule as long as the baseline score is (strictly) proper.  $\square$

The integrals in eq. (1) can be numerically approximated by truncating the domain to  $[u_{min}, y]$  and  $[y, u_{max}]$  such that

$$\begin{aligned} ACPS(P, y; c) &\approx \sum_{i=1}^N w_{2,i}^y (c^2 - P(u_{2,i}^y))^2 \left[ \frac{1}{(1-c)^2} \mathbb{I}(P(u_{2,i}^y) > c) + \frac{1}{c^2} \mathbb{I}(P(u_{2,i}^y) \leq c) \right] \\ &+ \sum_{i=1}^N w_{1,i}^y ((1-c)^2 - (1 - P(u_{1,i}^y))^2) \left[ \frac{1}{(1-c)^2} \mathbb{I}(P(u_{1,i}^y) > c) + \frac{1}{c^2} \mathbb{I}(P(u_{1,i}^y) \leq c) \right], \end{aligned} \quad (2)$$

where  $(w_{1,i}^y, u_{1,i}^y)_i$  and  $(w_{2,i}^y, u_{2,i}^y)_i$ , for  $i = 1, \dots, N$ , are the weights and locations of two Gaussian quadratures of  $N$  points on  $[y, u_{max}]$  and  $[u_{min}, y]$ , respectively.

*Remark 1.* In Bayesian statistics it is current practice the use of predictive distributions, mostly in the form of Monte Carlo samples from posterior predictive distributions of quantities of interest. The asymmetric scoring rule ACPS can be easily computed using the output of a Markov chain Monte Carlo algorithm by approximating the predictive distribution via the empirical cumulative distribution function (empirical CDF) and using it as a probabilistic forecast  $P$ .

To get an insight of the shape of the ACPS for varying levels of asymmetry,  $c$ , we consider two examples: one with several probabilistic forecasts and the other with a fixed forecast.




Figure 1: Asymmetric scoring rule  $ACPS(P, y; c)$  for different forecasting densities  $P$  and asymmetry level  $c$ . The observed value is fixed at  $y = 0$  and the true density is  $\mathcal{N}(0, 4)$ . Left panel: cumulative distribution functions of true density (solid, black) and forecasting densities:  $\mathcal{N}(-3, 1)$  (dashed, blue),  $\mathcal{N}(0, 1)$  (dashed, orange),  $\mathcal{N}(3, 1)$  (dashed, yellow),  $\mathcal{N}(0, 16)$  (dashed, purple). Right panel: value of the asymmetric scoring rule  $ACPS(P, y; c)$  against the asymmetry level  $c \in \{0.05, 0.275, 0.5, 0.725, 0.95\}$ , for each forecasting density (same colors as left panel).

**Example 1.** Let us consider several Gaussian probabilistic forecasts  $P$ . In Fig. 1 we show the value of the score on a range of asymmetry values  $c \in \{0.05, 0.275, 0.5, 0.725, 0.95\}$ , for a given observation  $y$  whose true density is a standard Gaussian. When the density forecast is Gaussian with the same mean as the target, the score is an inverse U-shaped function of the asymmetry level  $c$ . This is essentially due to the symmetry of the Gaussian distribution around its mean, since the probability mass in excess on the right tail is exactly equal to the mass lacking on the left one. However, notice that a higher score is assigned to  $\mathcal{N}(0, 1)$ , as compared to  $\mathcal{N}(0, 16)$ . Instead, the density forecasts  $\mathcal{N}(-3, 1)$  and  $\mathcal{N}(3, 1)$  receive a high penalty for high and small levels of  $c$ , respectively. This shows that values of  $c$  close to 1 heavily penalise forecasting densities that put more mass on the left part of the support as compared to the target, and conversely for values of  $c$  close to 0.



**Example 2.** Let us consider an alternative case when we keep fixed the probabilistic forecast to  $\mathcal{N}(2, 1)$  and inspect the value of the ACPS for alternative target densities. As expected (see Fig. 2), when the true density assigns more mass on the left part of the support as compared to the  $\mathcal{N}(2, 1)$ , the forecast receives a very low score especially for  $c$  close to 0. Conversely, when the underlying true density is  $\mathcal{N}(3, 1)$  the forecast receives a higher reward for  $c = 0.05$ , since its CDF is basically a left-shifted version of the target.




Figure 2: Asymmetric scoring rule  $ACPS(P, y; c)$  for different observed values  $y$  and asymmetry level  $c$ . The forecasting density is fixed at  $P = \mathcal{N}(2, 1)$ . Left panel: cumulative distribution functions of the forecasting density (dashed, black) and of observation densities:  $\mathcal{N}(3, 1)$  (solid, blue),  $\mathcal{N}(-1, 1)$  (solid, orange),  $\mathcal{N}(1, 4)$  (solid, yellow),  $\mathcal{N}(3, 4)$  (solid, purple). Right panel: value of the asymmetric scoring rule  $ACPS(P, y; c)$  against the asymmetry level  $c \in \{0.05, 0.275, 0.5, 0.725, 0.95\}$ , for each observation density (same colors as left panel).

## 2.1 Threshold and quantile-weighted versions

In addition to asymmetric preferences towards under- or overestimation, a decision maker is usually concerned with a precise forecast in a specific range of all possible values. Therefore, it is important to have a tool that allows to assign heterogeneous weights to various regions of the set of possible values of the variable. This calls for a scoring rule able to account for both error-weighting, i.e. asymmetric preferences and domain-weighting of density forecasts.

Gneiting and Ranjan (2011) modified the CRPS by re-weighting the loss according to a user-specified weight function, which allows to select the regions where the decision-maker has greater concern. By exploiting the representation of the CRPS in terms of quantile functions, they define a threshold-weighted (tCRPS) and quantile-weighted (qCRPS) score functions as follows

$$tCRPS(P, y) = \int_{-\infty}^{+\infty} |P(z) - \mathbb{I}(y \leq z)|^2 w(z) dz, \quad (3)$$

$$qCRPS(P, y) = \int_0^1 2(\mathbb{I}(y \leq P^{-1}(\alpha)) - \alpha)(P^{-1}(\alpha) - y)v(\alpha) d\alpha, \quad (4)$$

where  $w(z) \geq 0$  and  $v(\alpha) \geq 0$  are the weight functions and level  $\alpha \in (0, 1)$ . Table 1 reports some examples of weighting functions for the case of real-valued variables of interest; notice that the uniform weight,  $w(z) = 1$  and  $v(\alpha) = 1$ , leads to the standard CRPS. See Lerch et al. (2017) for discussion and applications of these scoring rules.

Table 1: Examples of weight functions for threshold-weighted and quantile-weighted CRPS, and variables supported on the real line.  $\phi, \Phi$  denote the probability density and cumulative distribution functions of the standard Normal distribution, respectively, with  $x \in \mathbb{R}$  and  $\alpha \in (0, 1)$ .

Emphasis	Threshold weight function	Quantile weight function
uniform	$w(x) = 1$	$v(\alpha) = 1$
center	$w(x) = \phi(x)$	$v(\alpha) = \alpha(1 - \alpha)$
tails	$w(x) = 1 - \phi(x)/\phi(0)$	$v(\alpha) = (2\alpha - 1)^2$
right tail	$w(x) = \Phi(x)$	$v(\alpha) = \alpha^2$
left tail	$w(x) = 1 - \Phi(x)$	$v(\alpha) = (1 - \alpha)^2$

The definition of ACPS in (1) can be modified to address this issue and obtain a threshold-weighted and a quantile-weighted asymmetric scoring rule, as follows.

**Definition 2** (Threshold-weighted ACPS). *Let  $G(du)$  be a positive measure<sup>1</sup>. We define the threshold-weighted asymmetric continuous probability score (tACPS), as*

$$\begin{aligned}
 tACPS(P, y; c) &= \int_{-\infty}^y (c^2 - P(u)^2) \left[ \frac{1}{(1-c)^2} \mathbb{I}(P(u) > c) + \frac{1}{c^2} \mathbb{I}(P(u) \leq c) \right] G(du) \\
 &+ \int_y^{+\infty} ((1-c)^2 - (1-P(u))^2) \left[ \frac{1}{(1-c)^2} \mathbb{I}(P(u) > c) + \frac{1}{c^2} \mathbb{I}(P(u) \leq c) \right] G(du),
 \end{aligned} \tag{5}$$

where  $c \in (0, 1)$  is the level of asymmetry and  $P$  is the probabilistic forecast and  $y$  the value that materializes.

**Definition 3** (Quantile-weighted ACPS). *Let  $p(u)$  denote the probability density function of  $P(u)$  and let  $P^{-1}(\alpha)$  be the corresponding quantile function at  $\alpha \in [0, 1]$ . Let  $V(d\alpha)$  be a positive measure on the unit interval. We define the quantile-weighted asymmetric continuous probability score (qACPS), as*

$$\begin{aligned}
 qACPS(P, y; c) &= \int_0^{P(y)} (c^2 - \alpha^2) \left[ \frac{1}{(1-c)^2} \mathbb{I}(\alpha > c) + \frac{1}{c^2} \mathbb{I}(\alpha \leq c) \right] \frac{1}{p(P^{-1}(\alpha))} V(d\alpha) \\
 &+ \int_{P(y)}^1 ((1-c)^2 - (1-\alpha)^2) \left[ \frac{1}{(1-c)^2} \mathbb{I}(\alpha > c) + \frac{1}{c^2} \mathbb{I}(\alpha \leq c) \right] \frac{1}{p(P^{-1}(\alpha))} V(d\alpha).
 \end{aligned} \tag{6}$$

As stated for ACPS, we can provide evidence of the properness of the two novel scores defined in eq. (5) and eq. (6).

**Theorem 2** (Properness of tACPS, qACPS). *For any  $c \in (0, 1)$ , it holds:*

- a) *the threshold-weighted asymmetric continuous probability score tACPS in eq. (5) is strictly proper;*
- b) *the quantile-weighted asymmetric continuous probability score qACPS in eq. (6) is strictly proper.*

*Proof.* The result follows from Theorem 1 and Matheson and Winkler (1976).  $\square$

Both tACPS and qACPS can be computed by approximating eq. (5) and eq. (6) in a way analogous to eq. (2). The main advantage of the tACPS and qACPS consists in the ability

<sup>1</sup>Notice that  $G(du)$  is not required to be a probability measure.

to consider two levels of asymmetry: in terms of the loss at each point, and over different regions of the domain. This is fundamental to answer the need of the decision maker who is concerned with the performance of the forecast in a given interval of possible values (e.g., the right tail) and who has an aversion to particular deviations from the target (e.g., averse to underestimation).

Tab. 2 provides a summary of some key differences between the CRPS and ACPS, and the corresponding weighted versions.

Table 2: Examples of scoring rules for evaluating density forecasts.

		Domain	
		uniform	weighted
Loss	symmetric	CRPS	tCRPS, qCRPS
	asymmetric	ACPS	tACPS, qACPS

*Remark 2* (Multivariate case). The proposed asymmetric scores can be easily generalized to multivariate settings. To this aim, denote with  $\mathcal{Q}$  the class of the Borel probability measures on  $\mathbb{R}^n$  and let  $F \in \mathcal{Q}$  be a probabilistic forecast identified via its cumulative distribution function,  $P$ . Let  $c \in (0, 1)$  represent the level of asymmetry and denote with  $\mathbf{y} = (y_1, \dots, y_n)'$  the multivariate value that materializes. The multivariate version of the asymmetric continuous probability score is defined as

$$\begin{aligned}
 ACPS(P, \mathbf{y}; c) &= \\
 &= \int_{-\infty}^{y_n} \cdots \int_{-\infty}^{y_1} (c^2 - P(\mathbf{u})^2) \left[ \frac{1}{(1-c)^2} \mathbb{I}(P(\mathbf{u}) > c) + \frac{1}{c^2} \mathbb{I}(P(\mathbf{u}) \leq c) \right] d\mathbf{u} \\
 &+ \int_{y_n}^{+\infty} \cdots \int_{y_1}^{+\infty} ((1-c)^2 - (1-P(\mathbf{u}))^2) \left[ \frac{1}{(1-c)^2} \mathbb{I}(P(\mathbf{u}) > c) + \frac{1}{c^2} \mathbb{I}(P(\mathbf{u}) \leq c) \right] d\mathbf{u},
 \end{aligned} \tag{7}$$

where  $d\mathbf{u} = du_1 \cdots du_n$ . Moreover, one can define a multivariate threshold-weighted ACPS by substituting the product of Lebesgue measures in eq. (7) with a positive measure  $G(d\mathbf{u})$  on  $\mathbb{R}^n$ .

In the literature on asymmetric point forecasting measures, the choice of the shape parameter(s) of the loss function has been of interest especially over the last decade. Some works, such as [Christoffersen and Diebold \(1996\)](#) and [Demetrescu and Hoke \(2019\)](#) performed an empirical exercise to rank competing forecasting models under asymmetric point forecast measure using a grid of asymmetry values.

Instead, [Elliott et al. \(2008, 2005\)](#) and [Patton and Timmermann \(2007\)](#) introduced procedures for inferring the value of the shape parameters of the forecaster's loss function. Assuming a collection of time series of forecasts of a given quantity of interest is available, they treat the loss function parameters as variable to be estimated and look for the values that would be most consistent with forecast rationality. This approach is appealing since the loss function parameters may provide information about the forecaster's objectives. However, the main drawback of these approaches is that they rely on the availability of a time series of

observed forecasts from relevant decision-makers (e.g., the IMF, the OECD, or central bankers). Therefore, the application is case specific for those data.

## 2.2 Testing predictive ability

When forecasts from multiple models are available, there is the need for statistical tools, such as tests, for assessing whether different forecasts are equally good. In the context of point forecasts, the Diebold-Mariano (DM) test is the most frequently used test for equal forecast performance. Essentially, it is based on the loss differential, defined as

$$d_t = L(e_{1,t}) - L(e_{2,t}),$$

where  $e_{j,t} = \hat{y}_{j,t} - y_t$  is the forecast error of model  $j = 1, 2$  at time  $t = 1, \dots, T$ ,  $\hat{y}_{j,t}$  is the point forecast of model  $j$ ,  $y_t$  is the true value, and  $L(\cdot)$  is a given loss function. The null hypothesis of equal accuracy in forecasting is  $H_0 : \mathbb{E}[d_t] = 0$  for all  $t$ , versus the alternative  $H_1 : \mathbb{E}[d_t] \neq 0$ . It can be shown that, if the loss differential series is (i) covariance stationary, and (ii) has short memory (see e.g. [McCracken \(2020\)](#)), then under the null hypothesis

$$\frac{\sqrt{T}\bar{d}}{\sqrt{2\pi f_d(0)}} \rightarrow \mathcal{N}(0, 1),$$

where  $\bar{d}$  and  $f_d(0)$  are the sample mean and the spectral density (at frequency 0) of the loss differential. Recently, [McCracken \(2020\)](#) found that the slow decay of the loss differential series is the most frequent problem that hampers the use of the Diebold-Mariano test in real data economic applications.

The density forecasting approach requires some adaptations of the Diebold-Mariano test, since the forecast is an infinite dimensional object  $P$ .

*Remark 3* (Modified DM test). To test the null hypothesis of equal accuracy of two competing models in a density forecasting approach, we modify the definition of the loss differential as follows. First, consider a proper scoring rule  $S$ , such as the ACPS or the CRPS, and denote the associated loss with  $S^*(y, P) = -S(y, P)$ . Then, the loss differential is defined as

$$d_t^* = S^*(y_t, P_{1,t}) - S^*(y_t, P_{2,t}). \quad (8)$$


Notice that the series  $d_t^*$  has the same interpretation as  $d_t$  in the original DM test, and following the same theoretical arguments one can prove that, under the null hypothesis  $H_0 : \mathbb{E}[d_t^*] = 0$  for each  $t$ , one has

$$\frac{\sqrt{T}\bar{d}^*}{\sqrt{2\pi f_{d^*}(0)}} \rightarrow \mathcal{N}(0, 1), \quad (9)$$

where  $\bar{d}^*$  and  $f_{d^*}(0)$  are the equivalent of  $\bar{d}$  and  $f_d(0)$  for  $d_t^*$ .

### 3 Illustrations and comparison with weighted CRPS


This section investigates the performance of the proposed asymmetric scoring rule and compares it with the CRPS. In order to assess the good performance of our measure, we consider different target densities: (i) Gaussian, (ii) Student-t, (iii) Gamma, (iv) Beta. This range includes families of distributions with different support ( $\mathbb{R}$ ,  $\mathbb{R}_+$  and  $[0, 1]$ ), skewed and with fat tails. For the asymmetric scoring rule ACPS we use varying levels of asymmetry, corresponding to  $c \in \{0.05, 0.275, 0.50, 0.725, 0.95\}$ . Recall that  $c = 0.50$  implies a symmetric loss.



	Forecasting density			
	$\mathcal{N}(0,1)$	$\mathcal{N}(-3,1)$	$\mathcal{N}(3,1)$	$\mathcal{N}(0,16)$
CRPS	1	4	3	2
ACPS( $\cdot, \cdot; 0.05$ )	1	2	4	3
ACPS( $\cdot, \cdot; 0.275$ )	1	3	4	2
ACPS( $\cdot, \cdot; 0.5$ )	1	4	3	2
ACPS( $\cdot, \cdot; 0.725$ )	1	4	3	2
ACPS( $\cdot, \cdot; 0.95$ )	1	4	2	3

Figure 3: Ranking of probabilistic forecasts. Results from  $S = 1$  simulation of  $N = 100$  observations. Density estimated with  $M = 500$  draws from forecasting distribution. Target is  $\mathcal{N}(0,1)$  (black), forecasting densities are:  $\mathcal{N}(0,1)$  (blue),  $\mathcal{N}(-3,1)$  (orange),  $\mathcal{N}(3,1)$  (yellow),  $\mathcal{N}(0,16)$  (purple).

Fig. 3 and Fig. 4 provide graphical evidence of the properness of the ACPS in two cases, with a Gaussian and a Student-t target, respectively. Both figures show that the ACPS rewards the forecast density which corresponds to the ground truth, for all levels of asymmetry. In addition, we find that the ranking of the competing probabilistic forecasts changes according to the value of  $c$ , due to the different penalty assigned to asymmetric deviations from the target.




	Forecasting density			
	$t(-3,1,3)$	$t(2,1,3)$	$t(0,1,5)$	$t(4,1,15)$
CRPS	3	2	1	4
ACPS( $\cdot, \cdot; 0.05$ )	3	2	1	4
ACPS( $\cdot, \cdot; 0.275$ )	3	2	1	4
ACPS( $\cdot, \cdot; 0.5$ )	3	2	1	4
ACPS( $\cdot, \cdot; 0.725$ )	4	2	1	3
ACPS( $\cdot, \cdot; 0.95$ )	4	2	1	3

Figure 4: Ranking of probabilistic forecasts. Results from  $S = 1$  simulation of  $N = 100$  observations. Density estimated with  $M = 500$  draws from forecasting distribution. Target is  $t(0,1,5)$  (black), forecasting densities are:  $t(-3,1,3)$  (blue),  $t(2,1,3)$  (orange),  $t(0,1,5)$  (yellow),  $t(4,1,15)$  (purple).

To investigate further this aspect, Fig. 5 presents the ranking of forecasts when none of the candidates corresponds to the true density, which is  $\mathcal{N}(2,4)$ . The CRPS indicates  $\mathcal{N}(3,1)$  as the best forecast, as does the ACPS for values of  $c$  around 0.5. However, when the ACPS assigns more weight to the asymmetric loss, that is for  $c$  close to the boundary of  $(0,1)$ , the ranking


significantly changes. For  $c = 0.05$ , that is when great importance is given to underestimation of the target, the  $\mathcal{N}(0, 1)$  is preferred, while  $\mathcal{N}(0, 16)$  is the best for the opposite case, when  $c = 0.95$ .



	Forecasting density			
	$\mathcal{N}(0, 1)$	$\mathcal{N}(-3, 1)$	$\mathcal{N}(3, 1)$	$\mathcal{N}(0, 16)$
CRPS	3	4	1	2
ACPS( $\cdot, \cdot; 0.05$ )	1	3	4	2
ACPS( $\cdot, \cdot; 0.275$ )	2	4	1	3
ACPS( $\cdot, \cdot; 0.5$ )	3	4	1	2
ACPS( $\cdot, \cdot; 0.725$ )	3	4	1	2
ACPS( $\cdot, \cdot; 0.95$ )	3	4	2	1

Figure 5: Ranking of probabilistic forecasts. Results from  $S = 1$  simulation of  $N = 100$  observations. Density estimated with  $M = 500$  draws from forecasting distribution. Target is  $\mathcal{N}(2, 4)$  (black), forecasting densities are:  $\mathcal{N}(0, 1)$  (blue),  $\mathcal{N}(-3, 1)$  (orange),  $\mathcal{N}(3, 1)$  (yellow),  $\mathcal{N}(0, 16)$  (purple).

Many economic and financial variables in levels are inherently positive (e.g. GDP, volatility) or take values on a bounded interval (e.g., interest rate, unemployment rate). To account for these cases, we investigate the performance of the ACPS in simulated experiments where the target density is either Gamma or Beta.



	Forecasting density			
	$\mathcal{Ga}(1, 1)$	$\mathcal{Ga}(2, 1)$	$\mathcal{Ga}(\frac{3}{2}, \frac{3}{2})$	$\mathcal{Ga}(1, 2)$
CRPS	3	1	4	2
ACPS( $\cdot, \cdot; 0.05$ )	4	1	2	3
ACPS( $\cdot, \cdot; 0.275$ )	4	1	2	3
ACPS( $\cdot, \cdot; 0.5$ )	4	1	2	3
ACPS( $\cdot, \cdot; 0.725$ )	4	1	2	3
ACPS( $\cdot, \cdot; 0.95$ )	2	1	3	4


Figure 6: Ranking of probabilistic forecasts. Results from  $S = 1$  simulation of  $N = 100$  observations. Density estimated with  $M = 500$  draws from forecasting distribution. Target is  $\mathcal{Ga}(2, 1)$  (black), forecasting densities are:  $\mathcal{Ga}(1, 1)$  (blue),  $\mathcal{Ga}(2, 1)$  (orange),  $\mathcal{Ga}(\frac{3}{2}, \frac{3}{2})$  (yellow),  $\mathcal{Ga}(1, 2)$  (purple).

Fig. 6 presents the results for a  $\mathcal{Ga}(2, 1)$  target density. By looking at the worst performing densities according to ACPS, we find that  $\mathcal{Ga}(1, 1)$  is assigned the highest penalty for values  $c \leq 0.725$ , while  $\mathcal{Ga}(1, 2)$  becomes the worst for  $c = 0.95$ . This reflects that for  $c \leq 0.725$ , the asymmetric score penalizes more the underestimation, while for  $c = 0.95$  it gives more weight to overestimation. Similar results are found in Fig. 7 with a positively skewed Beta target density,  $\mathcal{Be}(1, 2)$ .

### 3.1 Threshold-weighted version

We deep further the properties of the proposed asymmetric scoring rule by considering a threshold-weighted version and comparing it with the threshold-weighted CRPS. The goal is to disentangle the different role of the domain-weighting scheme, which reflects the interest





	Forecasting density			
	$\mathcal{Be}(1, 1)$	$\mathcal{Be}(1, 5)$	$\mathcal{Be}(1, 2)$	$\mathcal{Be}(5, 5)$
CRPS	3	2	1	4
ACPS( $\cdot, \cdot; 0.05$ )	2	3	1	4
ACPS( $\cdot, \cdot; 0.275$ )	3	2	1	4
ACPS( $\cdot, \cdot; 0.5$ )	2	3	1	4
ACPS( $\cdot, \cdot; 0.725$ )	3	4	1	2
ACPS( $\cdot, \cdot; 0.95$ )	2	4	1	3

Figure 7: Ranking of probabilistic forecasts. Results from  $S = 1$  simulation of  $N = 100$  observations. Density estimated with  $M = 500$  draws from forecasting distribution. Target is  $\mathcal{Be}(1, 2)$  (black), forecasting densities are:  $\mathcal{Be}(1, 1)$  (blue),  $\mathcal{Be}(1, 5)$  (orange),  $\mathcal{Be}(1, 2)$  (yellow),  $\mathcal{Be}(5, 5)$  (purple).

of the decision-maker in having good forecasts within a specific interval of values, and of the error-weighting scheme, which corresponds to the decision-maker’s loss in case of under or overestimation.

Consider a simulated experiment where  $N = 100$  observations are drawn from a Normal distribution  $\mathcal{N}(1, 4)$  and several forecasting densities are approximated using  $M = 500$  draws. We consider the domain-weighting schemes in Tab. 1, using 5 alternative asymmetry levels  $c \in \{0.05, 0.275, 0.50, 0.725, 0.95\}$ .

In Tab. 3 we find that the asymmetric penalty imposed by ACPS plays a significant role for all domain-weighting schemes considered. For an uniform weight, the ACPS agrees with the CRPS for  $c = 0.5$ , i.e. the symmetric case, but rewards differently the density forecasts for alternative values of the asymmetry level  $c$ . When the interest is focused on the right tail of the distribution, both threshold-weighted CRPS and ACPS agree, but when the attention is on the left tail, the two scoring rules perform remarkably different. The CRPS favours the standard Normal over the  $\mathcal{N}(3, 1)$ , while the ACPS rewards the latter for all  $c \geq 0.275$ .

The key insight obtained from this simulated exercise concerns the importance of domain- and error-weighting schemes. The first assigns an heterogeneous weight to the performance on different intervals, while the latter asymmetrically rewards negative and positive deviations from the true value. The threshold-weighted asymmetric scoring rule, tACPS, combines the two schemes and allows to evaluate the performance of the forecasting density from both perspectives. This is important to the decision makers, who are usually interested in a specific range of all possible values, thus calling for heterogeneous domain-weighting, and have asymmetric preferences towards under or overestimation, which motivates an asymmetric score.

Table 3: This table reports the ranking of probabilistic forecasts using  $t$ CRPS and  $t$ ACPS, for different weights (uniform, center, tails, right and left tail) and asymmetry levels ( $c \in \{0.05, 0.275, 0.5, 0.725, 0.95\}$ ). Results from  $S = 1$  simulation of  $N = 100$  observations (average score across all observations). Density estimated with  $M = 500$  draws from forecasting distribution. Target is  $\mathcal{N}(1, 4)$ , forecasting densities are  $\mathcal{N}(0, 1)$ ,  $\mathcal{N}(-3, 1)$ ,  $\mathcal{N}(3, 1)$ ,  $\mathcal{N}(0, 16)$ .

	$\mathcal{N}(0, 1)$	$\mathcal{N}(-3, 1)$	$\mathcal{N}(3, 1)$	$\mathcal{N}(0, 16)$
tCRPS uniform	4	2	3	1
tACPS( $\cdot, \cdot; 0.05$ ) uniform	1	3	4	2
tACPS( $\cdot, \cdot; 0.275$ ) uniform	2	1	4	3
tACPS( $\cdot, \cdot; 0.5$ ) uniform	4	2	3	1
tACPS( $\cdot, \cdot; 0.725$ ) uniform	4	3	2	1
tACPS( $\cdot, \cdot; 0.95$ ) uniform	4	3	1	2
tCRPS center	1	3	4	2
tACPS( $\cdot, \cdot; 0.05$ ) center	3	1	4	2
tACPS( $\cdot, \cdot; 0.275$ ) center	3	1	4	2
tACPS( $\cdot, \cdot; 0.5$ ) center	4	1	3	2
tACPS( $\cdot, \cdot; 0.725$ ) center	4	1	3	2
tACPS( $\cdot, \cdot; 0.95$ ) center	4	1	3	2
tCRPS tails	1	3	4	2
tACPS( $\cdot, \cdot; 0.05$ ) tails	1	4	2	3
tACPS( $\cdot, \cdot; 0.275$ ) tails	1	3	2	4
tACPS( $\cdot, \cdot; 0.5$ ) tails	2	3	4	1
tACPS( $\cdot, \cdot; 0.725$ ) tails	3	4	2	1
tACPS( $\cdot, \cdot; 0.95$ ) tails	4	3	1	2
tCRPS right tail	2	3	4	1
tACPS( $\cdot, \cdot; 0.05$ ) right tail	3	2	4	1
tACPS( $\cdot, \cdot; 0.275$ ) right tail	3	2	4	1
tACPS( $\cdot, \cdot; 0.5$ ) right tail	4	2	3	1
tACPS( $\cdot, \cdot; 0.725$ ) right tail	4	3	2	1
tACPS( $\cdot, \cdot; 0.95$ ) right tail	4	3	1	2
tCRPS left tail	1	3	2	4
tACPS( $\cdot, \cdot; 0.05$ ) left tail	1	3	4	2
tACPS( $\cdot, \cdot; 0.275$ ) left tail	2	3	1	4
tACPS( $\cdot, \cdot; 0.5$ ) left tail	4	3	1	2
tACPS( $\cdot, \cdot; 0.725$ ) left tail	4	3	1	2
tACPS( $\cdot, \cdot; 0.95$ ) left tail	4	2	1	3

## 4 Empirical applications

In the empirical applications, we adopt a similar framework to [Amisano and Giacomini \(2007\)](#) and [Gneiting and Ranjan \(2011\)](#), and consider the task of comparing density forecasts in a time series context. We use a fixed-length rolling window to provide a density forecast for  $h$  step ahead future observations. We focus on three different applications related to macroeconomics (employment growth rate) and to commodity prices (oil prices and electricity prices). We compare several univariate models, such as the autoregressive (AR) model, the Markov-switching (MS) AR model and the time-varying parameter (TVP) AR model.

We use the AR(1) as benchmark model, then we specify 12 lags for the employment

growth rate (i.e., 1 year of monthly observations) and 20 lags for the oil (i.e., 1 month of daily observations). Regarding the electricity prices, we include 7 lags (i.e., 1 week of daily observations) and by following common practice in the literature, we restrict lags to  $t - 1$ ,  $t - 2$ , and  $t - 7$ , which correspond to the previous day, two days before, and one week before the delivery time, recalling first similar conditions that may have characterized the market over the same hours and similar days (such as congestions and blackouts) and secondly the demand level during the days of the week. For the MS-AR model we consider only 1 lag, while for the TVP-AR model we use 1 and 2 lags. For both AR and TVP-AR, we consider three specifications of the variance: constant volatility and time-varying volatility in the form of stochastic volatility with Gaussian and Student-t error. For the MS-AR, we impose an identification constraint on the error variance.

In the first application, we aim at forecasting monthly US total nonfarm seasonally adjusted employment growth rate downloaded from the FRED database. We consider the growth rate of the monthly employment rate in US from January 1980 to April 2020. We see evidence of some spikes, in particular with a strong fall in April 2020 due to present COVID-19 situation (see Figure S.1 in the supplementary material). We use a rolling window approach of 20 years (thus 240 observations) and we forecast  $h = 1$  and  $h = 12$  (thus 1 year ahead) month ahead by using a recursive forecasting exercise.

For oil prices, we analyze daily West Texas Index (WTI) data (no weekends) from 02 January 2012 to 07 May 2020 in order to include in the analysis the recent turmoil. Indeed, large drops in demand that suddenly occurred and storage scarcity have resulted in negative WTI oil prices at the end of April 2020. As for the employment rate, we have used a rolling window of 4 years and we forecast  $h = 1$  and  $h = 5$  days ahead by applying recursive techniques.

In the third application, we consider the problem of forecasting the day-ahead electricity prices in Germany, one of the largest and leading energy market. In the electricity markets, the phenomenon of negative prices – when allowed to occur, such as in Germany where there is no floor price – has become more frequent due to the increasing share of electricity generated from renewable energy sources (RES) and the current impossibility to store it (see Figure 2 in the supplementary material). We analyze daily data (with weekends) from 01 January 2014 to 08 May 2020. For the forecasting analysis, we have considered a rolling window of 3 years and a recursive techniques for predicting  $h = 1$  and  $h = 7$  days ahead.

As we discussed in the introduction, policymakers or energy producers may be more concerned with forecasting values below a given threshold than the full distribution, since they require different measures, including in the case of energy variables to stop the production.<sup>2</sup> This supports the application of the ACPS. For the oil series we perform a case study around the collapse of WTI prices and discuss how the ACPS results can be applied to identify the true unknown density.

---

<sup>2</sup>Unfortunately, we have not precise data to compute (i) the value of this threshold, excluding the case of RES producers of electricity prices, that could be still profitable even when prices are marginally above zero, and (ii) the level of asymmetry of the loss function. Therefore, we investigate several values of  $c$ , the parameter that drives the asymmetry of our measure.

Before evaluating the relative performance of all models, we check the calibration of the density forecasts. Calibration of density forecasts is based on properties of a density and refers to absolute accuracy (see [Bassetti et al. \(2019\)](#) for further details). The absolute accuracy can be studied by testing forecast accuracy relative to the “true”, unobserved density. [Dawid \(1982\)](#) introduced the criterion of calibration for comparing prequential probabilities with binary random outcomes and exploited the concept of probability integral transform (PIT), that is the value that a predictive CDF attains at the observations, for continuous random variables. The PITs summarize the properties of the densities and may help us to judge whether the densities are biased in a particular direction and whether the width of the densities has been roughly correct on average, see [Diebold et al. \(1998\)](#). The PITs can provide an indication of whether a density is wrong in predicting higher moments or specific parts of the distribution, such as the tails; however they cannot distinguish among models that are also correctly calibrated. We apply the test of [Knuppel \(2015\)](#) and refer to [Rossi and Sekhposyan \(2013\)](#) for evaluation of PITs in presence of instabilities, [Rossi and Sekhposyan \(2014\)](#) for application with large database and [Rossi and Sekhposyan \(2019\)](#) for a comparison of alternative tests for correct specification of density forecasts.

The PIT tests in [Tab. 4](#) indicate that all densities are correctly calibrated for the employment growth rate at 5% significance level, excluding the one given by the TVP-AR(2) model at 12-month horizon, for which the p-value is marginally lower at 4.9%. Density forecasts from models TVP-AR(2)-SV<sup>3</sup> and TVP-AR(2)-tSV are calibrated at 1-day ahead horizon; no density is correctly calibrated at 5-days ahead horizons. All densities are not correctly calibrated when predicting EEX electricity prices at both horizons. So, the PITs analysis suggests there is not a stochastically dominating model, but more specifications can provide (absolute) accurate forecasts suggesting the use of relative metrics such as the ACPStest to discriminate among them. In the case of EEX prices, all models are wrong and a possible explanation is that the models considered in this text are based only on econometric properties of the series, hence they may be labelled as “purely econometric” models. [Gianfreda et al. \(2020a\)](#) and [Gianfreda et al. \(2020b\)](#) document how important is to extend these models with economically relevant variables, such as variables related to the demand and the production of electricity, including renewable energy sources, to increase accuracy. We leave this extension for further research and apply our metrics to an example where models in terms of calibration are all wrong.

[Tab. 4](#) shows the ranking of the probability forecasts over vintages and across models for all the three datasets for  $c = 0.05, 0.5, 0.95$ .<sup>4</sup> The ACPStest presented in [Section 2.2](#) is also reported.<sup>5</sup>

Regarding the employment growth rate, we can see at horizon 1-month ahead that the best model for  $c = 0.05$  is the AR(12)-tSV, for  $c = 0.5$  it is the TVP-AR with 2 lags (the

<sup>3</sup>Notice that the TVP-AR(2)-SV is always preferred in terms of relative accuracy.

<sup>4</sup>See [Table IV](#) in the supplementary material for results for a higher range of  $c$ .

<sup>5</sup>In order to perform the test, we checked the stationarity and short memory of the loss differential series using the ADF test and the autocorrelation function, respectively.

Table 4: Ranking of probability forecasts and accuracy test. Best model, over vintages, according to: CRPS; ACPS with  $c = 0.05; 0.5; 0.95$  for the three different datasets: Employment (top); Oil (middle) and EEX (bottom).

EMPL													
Horizon 1	AR(1)	AR(1)-SV	AR(1)-tSV	AR(12)	AR(12)-SV	AR(12)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
ACPS(.,.;0.05)	12	9***	5***	13	2***	1***	11***	3***	10***	8***	4***	7***	6***
ACPS(.,.;0.5)	13	10***	9***	11	8***	7***	12***	4***	6***	5***	1***	2***	3***
ACPS(.,.;0.95)	13	4***	3***	12	1***	2***	11***	5***	7***	9***	6***	8***	10***
CRPS	13	10***	9***	11	8***	7***	12***	4***	6***	5***	1***	2***	3***
Horizon 12	AR(1)	AR(1)-SV	AR(1)-tSV	AR(12)	AR(12)-SV	AR(12)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
ACPS(.,.;0.05)	11	12	10	13	3	2	9***	4	6	5	1	8	7
ACPS(.,.;0.5)	12	4***	3***	13	1***	2***	11***	5***	7***	8***	6***	9***	10***
ACPS(.,.;0.95)	12	1***	2***	13	3***	4***	11***	5***	7***	8***	6***	9***	10***
CRPS	12	4***	3***	13	1***	2***	11***	5***	7***	8***	6***	9***	10***
OIL													
Horizon 1	AR(1)	AR(1)-SV	AR(1)-tSV	AR(20)	AR(20)-SV	AR(20)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
ACPS(.,.;0.05)	12	8*	11	10**	7*	9**	13	6***	3***	4***	5***	2***	1***
ACPS(.,.;0.5)	12	9	13	7	8	11	10	3***	6***	5***	1***	4***	2***
ACPS(.,.;0.95)	12	8	11**	9	7	10**	13	3**	6*	5*	1**	4*	2*
CRPS	12	9	13	7	8	11	10	3***	6***	5***	1***	4***	2***
Horizon 5	AR(1)	AR(1)-SV	AR(1)-tSV	AR(20)	AR(20)-SV	AR(20)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
ACPS(.,.;0.05)	10	5	7	9	3	6	12	13	2	8	11	1*	4*
ACPS(.,.;0.5)	3	7	6	8	4	5	11	1**	13	12	2	10	9
ACPS(.,.;0.95)	6	1	3*	8	2	4	13	5	10	7	11	12	9
CRPS	3	6	7	8	4	5	11	1**	13	12	2	10	9
EEX													
Horizon 1	AR(1)	AR(1)-SV	AR(1)-tSV	AR(20)	AR(20)-SV	AR(20)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
ACPS(.,.;0.05)	11	8*	6***	9***	3***	1***	12	13	7**	4**	10	5**	2***
ACPS(.,.;0.5)	12	10***	11*	7***	2***	5***	13	9***	6***	4***	8***	3***	1***
ACPS(.,.;0.95)	12	7***	8***	11***	2***	6***	13	9***	3***	5***	10***	1***	4***
CRPS	13	10***	11*	7***	1***	5***	12	9***	6***	4***	8***	3***	2***
Horizon 7	AR(1)	AR(1)-SV	AR(1)-tSV	AR(20)	AR(20)-SV	AR(20)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
ACPS(.,.;0.05)	5	13	12	1***	11	4	7	3	9	8	2	10	6
ACPS(.,.;0.5)	10	12	13	9***	8***	7***	11	6***	4***	2***	5***	3***	1***
ACPS(.,.;0.95)	12	11	10	9***	7*	5**	13	1***	8*	4**	2***	6*	3**
CRPS	10	12	13	9***	8***	7***	11	6***	4***	2***	5***	3***	1***

Notes:  
<sup>1</sup>\*\*\*, \*\* and \* indicate scores are significantly different from 1 at 1%, 5% and 10%, according to the ACPStest in Section 2.2.  
<sup>2</sup> Gray cells indicate models that are correctly calibrated at 5% significance level according to the Knuppel test.

same for the CRPS measure), and for  $c = 0.95$  it is the AR(12)-SV, showing differences across different levels of asymmetry. The test indicates that most of the models provide superior forecasts than the AR(1) benchmark and only the AR(12) model does not provide gains. The difference in model performance for various levels of  $c$  is confirmed for  $h = 12$  and interesting for  $c = 0.05$  only the AR(1)-MS is statistically superior. Therefore, our evidence supports the large literature on the use of time-varying and nonlinear models in modelling and forecasting (un)employment data. Moreover, the best model for  $h = 12$  and  $c = 0.5$  is the same when applying the CRPS. In Fig. 8, we report the best model in each vintage for the two horizon ahead, where the black line refers to the CRPS, the red and the yellow for the ACPS for  $c = 0.05$  and for  $c = 0.95$ , respectively. The graph shows large instability in the best model, in particular when using the CRPS. The ACPS rules seem to prefer one of the alternative models for more consecutive vintages. For example, by looking at the relative frequency of occurrence of each model as the best model, we find that for  $c = 0.05$ , 31% times the AR(12)-tSV is considered the best model for  $h = 1$ . Similar percentages are found for other levels of  $c$  and  $h$ , despite model order varies substantially across measures.

Moving to the oil prices, in the middle panel of Tab. 4, we find that across vintages, for 1 day ahead the TVP-AR(2) is the best model whereas the TVP-AR(2)-SV is the second




Figure 8: Best model in each vintage for EMPL dataset: CRPS (black), ACPS with  $c = 0.05$  (red), ACPS with  $c = 0.95$  (yellow).

best for the asymmetric levels  $c = 0.5, 0.95$  and the CRPS. For  $c = 0.05$  the best model is the TVP-AR(2)-SV model, supporting PITS evidence that this model is among the few ones correctly calibrated. The TVP-AR(2)-SV model is again the best model for 1 week ahead of forecasting and for  $c = 0.05$  and it is one of the two models to be statistically superior to the AR benchmark. For the same weekly horizon and other levels of  $c$ , again only few models are superior to the benchmark. Fig. 9 confirms that the ACPS is less variable in this selection than the CRPS.

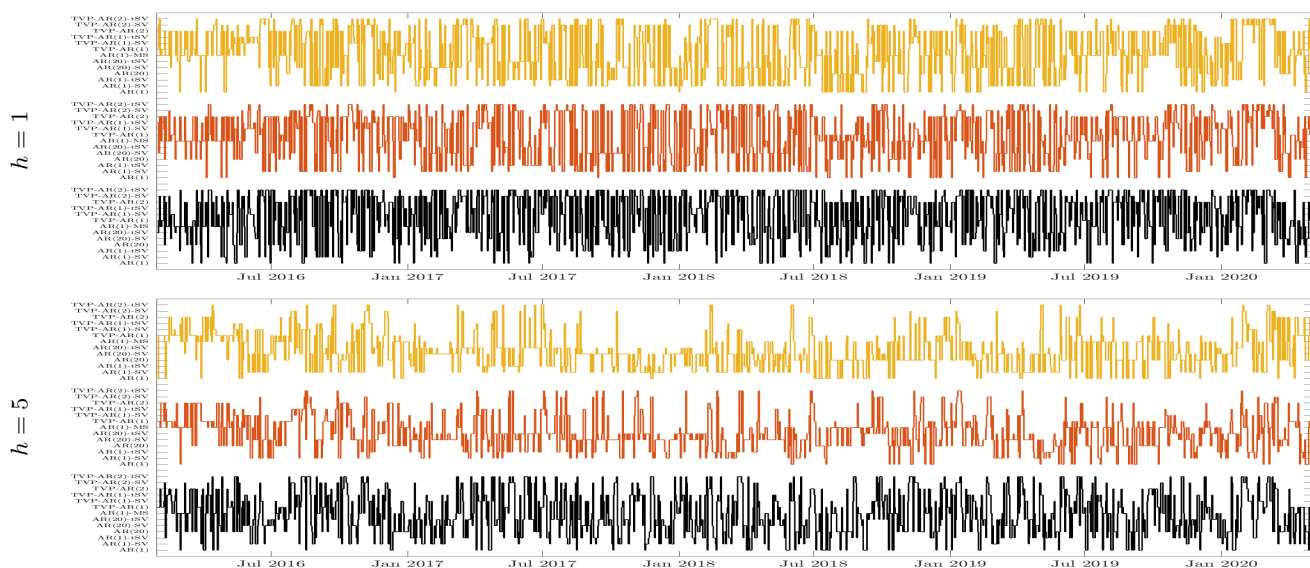


Figure 9: Best model in each vintage for OIL dataset: CRPS (black), ACPS with  $c = 0.05$  (red), ACPS with  $c = 0.95$  (yellow).

Fig. 10 illustrates the ACPS for one step ahead density forecasts of the OIL prices, according



to an AR(20) model, for each vintage of the rolling estimation and various levels of asymmetry. Despite showing the results for a single model, this figure presents some interesting insights. By looking at the scores between April 17 and April 21, we find that the forecast is worst performing for  $c = 0.05$  and best for  $c = 0.95$ , indicating that the density forecast assigns more mass on the right part of the support as compared to the density of the observations. This situation is similar to the yellow line in Fig. 1. Surprisingly, the ranking is reversed between April 21 and April 24, where the forecast receives a higher score under  $c = 0.05$ . This suggests that the density forecast is likely to be a right-shifted version of the observation density, similarly to the blue line in Fig. 1.

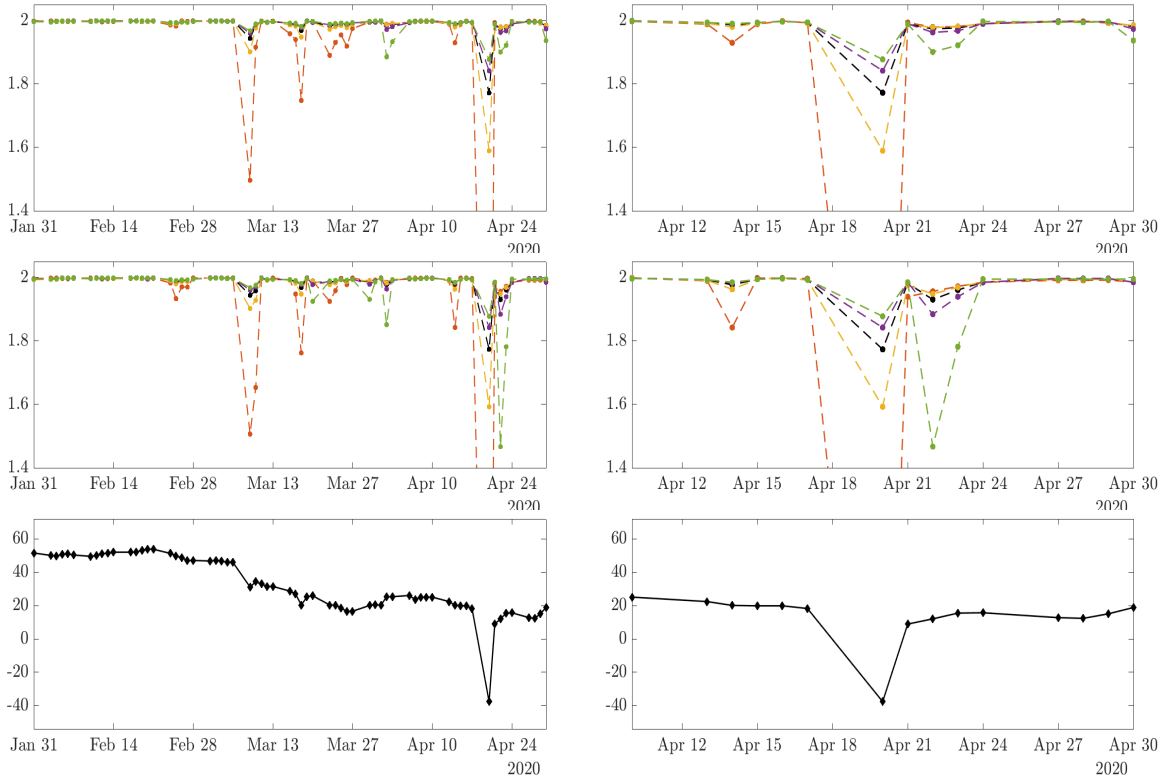


Figure 10: Top two rows: values of ACPS for one step ahead density forecasts of the OIL prices according to a TVP-AR(2) model and an AR(20) model, respectively, for selected rolling windows (x-axis) and different asymmetry levels  $c$ : 0.05 (dashed red line), 0.275 (dashed yellow line), 0.50 (dashed black line), 0.725 (dashed purple line), 0.95 (dashed green line). Bottom row: observed values of the time series (solid black line). The right column is a zoomed-in version of the left column.

These results highlight how accounting for asymmetry in forecast evaluation may lead to dramatically different implications. By looking at the period until April 21, a decision maker averse to overestimation of oil price is likely to discard the AR(20) in favor of alternative models for making forecasts. Conversely, another agent facing the same decision problem, equipped with the same data and models, but averse to underestimation, is likely to agree to the AR(20).

Moreover, these insights provide an important value added of the ACPS as compared to symmetric scores. By looking at variation of the ranking according to the ACPS over time, it is possible to infer the relative dynamics of the forecasting and observation densities. In the case previously mentioned, between April 17 and April 21 the forecast tends to overestimate

(i.e., its CDF is to the right of the observations CDF), while it tends to underestimate between April 21 and April 24 (i.e., its CDF is to the left of the observations CDF). Under a symmetric score it is not possible to grasp these insights since negative and positive deviations from the target are equally penalized.

The bottom panel of Tab. 4 reports the results for the electricity prices. As in the previous cases, there is large uncertainty on the model ranking. In line with PIT evidence, the high volatility, spikes and negative prices of the electricity prices drive different results depending on the level of asymmetry of the user. At  $h = 1$  and  $c = 0.05$ , the AR(20)-tSV is the best model, for higher values of  $c$ , the TVP-AR(2)-SV and TVP-AR(2)-tSV are the preferred ones. Many models with time-varying volatility outperform the constant volatility models, confirming evidence in [Gianfreda et al. \(2020b\)](#). At  $h = 12$  the AR(20) for  $c = 0.05$ , the TVP-AR(2)-tSV for  $c = 0.5$ , and the TVP-AR(1) for  $c = 0.95$  give the highest ACPS. Fig. 11 again indicates more stable performance of some models when accounting for asymmetry relative to use the symmetric CRPS.

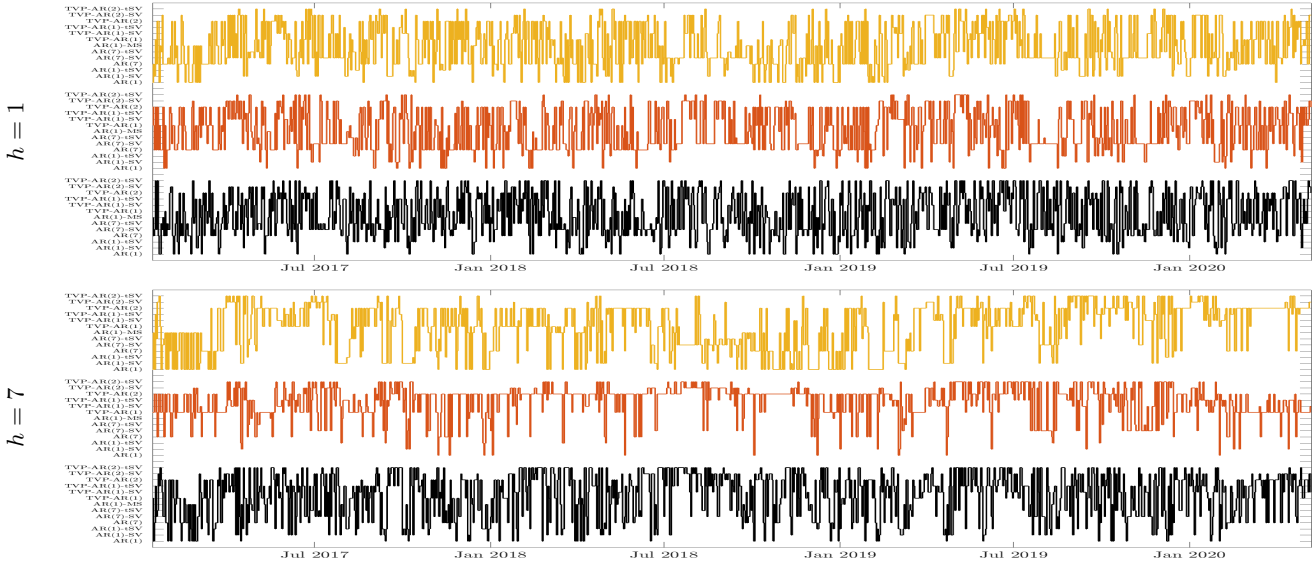


Figure 11: Best model in each vintage for EEX dataset: CRPS (black), ACPS with  $c = 0.05$  (red), ACPS with  $c = 0.95$  (yellow).

## 5 Conclusions

This paper has introduced a novel asymmetric proper score for probabilistic forecasts of continuous variables, the ACPS. Its main application is the evaluation and comparison of density forecasts. In addition, we have proposed a threshold- and quantile-weighted version of the asymmetric score, which, by reweighing the domain, allows for a further level of asymmetry in the evaluation of forecasts. We also derive a test to compare the statistical accuracy of different forecasts. The definition of ACPS is sufficiently flexible to be used in a variety of univariate contexts, and carries over to the multivariate case. The latter deserves further

investigation and is an open field for future research.

We provide a tool able to account for the decision maker's preferences in the evaluation of density forecasts both in terms of domain- and error-weighting schemes.

In an artificial data exercise, we have shown the good performance of our proposed asymmetric score for different continuous target distributions. In relevant macroeconomic and energy applications, we evaluate our score across different models and for different horizons, and we improve on the quality of the forecasts by providing an effective tool for density forecast evaluation.

The proposed score, ACPS, is of general use in any situation where the decision maker has asymmetric preferences in the evaluation of forecasts and thus it can be applied to a much wide range of applications. Further extensions could cover the area of forecast instability, see [Giacomini and Rossi \(2010\)](#), and the case of state-dependent function of economic variables such as in [Odendahl et al. \(2020\)](#).

## References

- Amisano, G. and R. Giacomini (2007). Comparing density forecasts via weighted likelihood ratio tests. *Journal of Business & Economic Statistics* 25(2), 177–190.
- Anscombe, F. (1968). Topics in the investigation of linear relations fitted by the method of least squares. *Journal of the Royal Statistical Society (Series B)* 29, 1–52.
- Artis, M. and M. Marcellino (2001). Fiscal forecasting: The track record of the IMF, OECD and EC. *The Econometrics Journal* 4(1), 20–36.
- Bassetti, F., R. Casarin, and F. Ravazzolo (2019). Density forecasting. In P. Fuleky (Ed.), *Macroeconomic Forecasting in the Era of Big Data*, Chapter 15, pp. 465–494. Springer.
- Bauer, H. (2011). *Measure and integration theory*, Volume 26. Walter de Gruyter.
- Boero, G., J. Smith, and K. F. Wallis (2008). Evaluating a three-dimensional panel of point forecasts: the Bank of England Survey of external forecasters. *International Journal of Forecasting* 24(3), 354–367.
- Carriero, A., T. E. Clark, and M. Marcellino (2020). Nowcasting tail risks to economic activity with many indicators. Technical report, Federal Reserve Bank of Cleveland, Working Paper No. 20-13.
- Carter, C. K. and R. Kohn (1994). On Gibbs sampling for state space models. *Biometrika* 81(3), 541–553.
- Christodoulakis, G. A. and E. C. Mamatzakis (2008). An assessment of the EU growth forecasts under asymmetric preferences. *Journal of Forecasting* 27(6), 483–492.

- Christodoulakis, G. A. and E. C. Mamatzakis (2009). Assessing the prudence of economic forecasts in the EU. *Journal of Applied Econometrics* 24(4), 583–606.
- Christoffersen, P. F. and F. X. Diebold (1996). Further results on forecasting and model selection under asymmetric loss. *Journal of Applied Econometrics* 11(5), 561–571.
- Christoffersen, P. F. and F. X. Diebold (1997). Optimal prediction under asymmetric loss. *Econometric Theory* 13(6), 808–817.
- Dawid, A. P. (1982). *Intersubjective statistical models*, pp. 217–232. North-Holland Publishing Company.
- Demetrescu, M. and S. H. Hoke (2019). Predictive regressions under asymmetric loss: Factor augmentation and model selection. *International Journal of Forecasting* 35(1), 80–99.
- Diebold, F. X., T. A. Gunther, and A. S. Tay (1998). Evaluating density forecasts with applications to financial risk management. *International Economic Review* 39(4), 863–83.
- Dovern, J. and N. Janssen (2017). Systematic errors in growth expectations over the business cycle. *International Journal of Forecasting* 33(4), 760–769.
- Elliott, G., I. Komunjer, and A. Timmermann (2008). Biases in macroeconomic forecasts: irrationality or asymmetric loss? *Journal of the European Economic Association* 6(1), 122–157.
- Elliott, G. and A. Timmermann (2004). Optimal forecast combinations under general loss functions and forecast error distributions. *Journal of Econometrics* 122(1), 47–79.
- Elliott, G. and A. Timmermann (2016a). *Economic Forecasting*. Princeton University Press.
- Elliott, G. and A. Timmermann (2016b). Forecasting in economics and finance. *Annual Review of Economics* 8, 81–110.
- Elliott, G., A. Timmermann, and I. Komunjer (2005). Estimation and testing of forecast rationality under flexible loss. *The Review of Economic Studies* 72(4), 1107–1125.
- Frühwirth-Schnatter, S. (2006). *Finite mixture and Markov switching models*. Springer Science & Business Media.
- Giacomini, R. and B. Rossi (2010). Forecast comparisons in unstable environments. *Journal of Applied Econometrics* 25(4), 595–620.
- Gianfreda, A., F. Ravazzolo, and L. Rossini (2020a). Comparing the forecasting performances of linear models for electricity prices with high RES penetration. *International Journal of Forecasting* 36(3), 974–986.
- Gianfreda, A., F. Ravazzolo, and L. Rossini (2020b). Large time-varying volatility models for electricity prices. Technical report, CAMP Working Paper Series 05/2020.

- Gneiting, T. (2011). Making and evaluating point forecasts. *Journal of the American Statistical Association* 106(494), 746–762.
- Gneiting, T. and A. E. Raftery (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102(477), 359–378.
- Gneiting, T. and R. Ranjan (2011). Comparing density forecasts using threshold- and quantile-weighted scoring rules. *Journal of Business & Economic Statistics* 29(3), 411–422.
- Gneiting, T. and R. Ranjan (2013). Combining predictive distributions. *Electronic Journal of Statistics* 7, 1747–1782.
- Granger, C. W. J. and M. H. Pesaran (2000). Economic and statistical measures of forecast accuracy. *Journal of Forecasting* 19, 537–560.
- Knuppel, M. (2015). Evaluating the calibration of multi-step-ahead density forecasts using raw moments. *Journal of Business & Economic Statistics* 33(2), 270–281.
- Lerch, S., T. Thorarinsdottir, F. Ravazzolo, and T. Gneiting (2017). Forecaster’s dilemma: Extreme events and forecast evaluation. *Statistical Science* 32(1), 106–127.
- Matheson, J. E. and R. L. Winkler (1976). Scoring rules for continuous probability distributions. *Management Science* 22(10), 1087–1096.
- McCracken, M. W. (2020). Diverging tests of equal predictive ability. *Econometrica* 88(4), 1753–1754.
- Odendahl, F., B. Rossi, and T. Sekhposyan (2020). Comparing forecast performance with state dependence. Technical report, Working Paper.
- Patton, A. J. and A. Timmermann (2007). Testing forecast optimality under unknown loss. *Journal of the American Statistical Association* 102(480), 1172–1184.
- Rossi, B. and T. Sekhposyan (2013). Conditional predictive density evaluation in the presence of instabilities. *Journal of Econometrics* 177(2), 199–212.
- Rossi, B. and T. Sekhposyan (2014). Evaluating predictive densities of us output growth and inflation in a large macroeconomic data set. *International Journal of Forecasting* 30(3), 662–682.
- Rossi, B. and T. Sekhposyan (2019). Alternative tests for correct specification of conditional predictive densities. *Journal of Econometrics* 208(2), 638–657.
- Savage, L. J. (1971). Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* 66(336), 783–801.
- Schervish, M. J. (1989). A general method for comparing probability assessors. *The Annals of Statistics* 17(4), 1856–1879.

- Timmermann, A. (2006). *Forecast Combinations*, Volume 1, Chapter 4, pp. 135–196. Elsevier.
- Tsuchiya, Y. (2016). Assessing macroeconomic forecasts for Japan under an asymmetric loss function. *International Journal of Forecasting* 32(2), 233–242.
- Winkler, R. L. (1994). Evaluating probabilities: Asymmetric scoring rules. *Management Science* 40(11), 1395–1405.
- Zarnowitz, V. (1969). Topics in the investigation of linear relations fitted by the method of least squares. *American Statistician* 23, 12–16.



# Supplementary Material to “Proper scoring rules for evaluating asymmetry in density forecasting”

## Abstract

This appendix contains the details of the model specifications used in the empirical applications in Section S.1. The data used in the empirical application is described in Section S.2, while Section S.3 provides additional details of the empirical analysis.

## S.1 Models

Let  $i \in \{EMPL, OIL, EEX\}$ .

### AR(1)

As a benchmark, we consider a standard autoregressive model of order 1, AR(1), defined as

$$y_{i,t} = \alpha_i + \beta_i y_{i,t-1} + \epsilon_{i,t}, \quad \epsilon_{i,t} \sim \mathcal{N}(0, \sigma_i^2), \quad (\text{S.1})$$

where  $t = 1, \dots, T$ . As for the other models, we also consider different lag specification.

### Markov Switching AR(1)

The assumption of constant parameters in time series analysis is likely to be incorrect, especially for long time series. Motivated by this fact, we estimate a Markov switching AR(1) model, called MS-AR(1), where the autoregressive parameter jumps according to a latent Markov chain

$$y_{i,t} = \alpha_{i,S_t} + \beta_{i,S_t} y_{i,t-1} + \epsilon_{i,t}, \quad \epsilon_{i,t} \sim \mathcal{N}(0, \sigma_{i,S_t}^2) \quad (\text{S.2})$$

where each  $S_{i,t}$  follows a homogeneous Markov chain with transition matrix

$$\Xi_i = \begin{bmatrix} \xi_{i,1,1} & \xi_{i,1,2} \\ \xi_{i,2,1} & \xi_{i,2,2} \end{bmatrix}$$

with  $\xi_{i,k,j} = P(S_{i,t} = j | S_{i,t} = k)$ .

### Bayesian inference

We sample the path of  $S_{i,t}$  using the forward filter, backward sampler of [Frühwirth-Schnatter \(2006\)](#).

Fix  $i$  and define  $\beta_{S_t} = (\alpha_i, \beta_{i,S_t})'$  and  $X_t = (1, y_{t-1})$ . Finally, for each state  $m = 1, \dots, M$  let  $\mathcal{T}_m = \{t = 1, \dots, T : s_t = m\}$ . Assume a Gaussian prior for  $\beta_m$ , for each state  $m = 1, \dots, M$ , as

$$P(\beta_m) = \mathcal{N}(\underline{\mu}_\beta, \underline{V}_\beta).$$

Then, the posterior is obtained as

$$\boldsymbol{\beta}_m | \mathbf{y}, \mathbf{s}_t, \sigma^2 \sim \mathcal{N}(\bar{\boldsymbol{\mu}}_\beta, \bar{\mathbf{V}}_\beta),$$

where

$$\bar{\mathbf{V}}_\beta = \left( \mathbf{V}_\beta^{-1} + \sum_{t \in \mathcal{T}_m} \frac{\mathbf{X}'_t \mathbf{X}_t}{\sigma_m^2} \right)^{-1}, \quad \bar{\boldsymbol{\mu}}_\beta = \bar{\mathbf{V}}_\beta \left( \mathbf{V}_\beta^{-1} \boldsymbol{\mu}_\beta + \sum_{t \in \mathcal{T}_m} \frac{\mathbf{X}'_t \mathbf{y}_t}{\sigma_m^2} \right).$$

Assume a Dirichlet prior for  $\boldsymbol{\xi}_m = (\xi_{m,1}, \dots, \xi_{m,m})$ , for each state  $m = 1, \dots, M$ , as

$$P(\boldsymbol{\xi}_m) = \text{Dir}(\mathbf{c}_m).$$

Then, the posterior is obtained as

$$\boldsymbol{\xi}_m | \mathbf{s}_t \sim \text{Dir}(\mathbf{c}_m + \mathbf{N}_m),$$

where

$$N_{m,l} = \#\{t : S_{t-1} = m \wedge S_t = l\}.$$

For identifying the states, we impose a restriction on the levels of the volatility, i.e.  $\sigma_{i,1}^2 < \sigma_{i,2}^2$ . Therefore, we label state 1 the low volatility regime and state 2 the high volatility regime.

### TVP-AR(1)

In order to capture potential smooth variations of the coefficient, we consider a TVP framework, denoted TVP-AR(1), where the latter parameter is assumed to evolve according to a latent AR(1) process

$$y_{i,t} = \alpha_{i,t} + \beta_{i,t} y_{i,t-1} + \epsilon_{i,t}, \quad \epsilon_{i,t} \sim \mathcal{N}(0, \sigma_i^2) \quad (\text{S.3})$$

$$\beta_{i,t} = A_i \beta_{i,t-1} + \boldsymbol{\eta}_{i,t}, \quad \boldsymbol{\eta}_{i,t} \sim \mathcal{N}(0, \Omega_i) \quad (\text{S.4})$$

We sample the path of  $\beta_{i,t}$  using the forward filter, backward sampler of [Carter and Kohn \(1994\)](#).

We assume a Gaussian prior for each column  $j$  of  $A_i$ , that is

$$P(A_{i,:j}) = \mathcal{N}(\underline{\boldsymbol{\mu}}_A, \underline{\mathbf{V}}_A),$$

thus obtaining a Gaussian posterior

$$A_{i,:j} | \{\boldsymbol{\beta}_{i,1}, \dots, \boldsymbol{\beta}_{i,T}\}, \Omega_i \sim \mathcal{N}(\bar{\boldsymbol{\mu}}_A, \bar{\mathbf{V}}_A),$$

where

$$\bar{V}_A = \left( \underline{V}_A^{-1} + \sum_{t=2}^T \beta'_{i,t-1} \Omega_i^{-1} \beta_{i,t-1} \right)^{-1}, \quad \bar{\mu}_A = \bar{V}_A \left( \underline{V}_A^{-1} \underline{\mu}_A + \sum_{t=2}^T \beta'_{i,t-1} \Omega_i^{-1} \beta_{i,t} \right).$$

### Stochastic Volatility

We also study the previous autoregressive models with the inclusion of stochastic volatility. We present only the extension for the AR(1) in eq. (S.1), and analogous changes apply to model (S.4).

Let  $h_{i,t}$ , for  $t = 1, \dots, T$ , be the log-variance for series  $i$  at time  $t$ , then an AR(1) model with Gaussian noise and stochastic volatility, denoted AR(1)-SV, is given by

$$\begin{aligned} y_{i,t} &= \alpha_i + \beta_i y_{i,t-1} + \epsilon_{i,t}, & \epsilon_{i,t} &\sim \mathcal{N}(0, \exp(h_{i,t})) \\ h_{i,t} &= h_{i,t-1} + u_{i,t}, & u_{i,t} &\sim \mathcal{N}(0, \sigma_{i,h}^2). \end{aligned} \quad (\text{S.5})$$

The last model considered is an extension of the AR(1) with stochastic volatility that assumes Student- $t$  noise. Let  $\nu_i > 0$  be the degrees of freedom parameter. An AR(1) with Student- $t$  noise and stochastic volatility, denoted AR(1)-SV- $t$ , is given by


$$\begin{aligned} y_{i,t} &= \alpha_i + \beta_i y_{i,t-1} + \epsilon_{i,t}, & \epsilon_{i,t} &\sim \mathcal{N}(0, \exp(h_{i,t})/\lambda_{i,t}) \\ h_{i,t} &= h_{i,t-1} + u_{i,t}, & u_{i,t} &\sim \mathcal{N}(0, \sigma_{i,h}^2), \end{aligned} \quad (\text{S.6})$$

where  $\lambda_{i,t} \sim \mathcal{IG}(\nu_i/2, \nu_i/2)$ .

## S.2 Data description

Figure S.1: **Raw data series.**

This figure reports the raw time series for EMPL, OIL, and EEX .



### S.3 Additional details on empirical applications

Table S.1: Relative frequency of occurrence of each model for EMPL as best model, across vintages, according to: CRPS (rows 1-2), ACPS with  $c = 0.05$  (rows 3-4), ACPS with  $c = 0.95$  (rows 5-6)

		AR(1)	AR(1)-SV	AR(1)-tSV	AR(12)	AR(12)-SV	AR(12)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
CRPS	$h = 1$	0.01	0.10	0.10	0.03	0.16	0.13	0.00	0.06	0.07	0.07	0.10	0.09	0.06
	$h = 12$	0.01	0.10	0.14	0.05	0.15	0.15	0.00	0.17	0.06	0.05	0.04	0.05	0.04
ACPS 0.05	$h = 1$	0.00	0.20	0.05	0.00	0.31	0.06	0.01	0.10	0.04	0.03	0.11	0.05	0.03
	$h = 12$	0.01	0.18	0.03	0.00	0.31	0.07	0.03	0.31	0.00	0.00	0.04	0.00	0.00
ACPS 0.95	$h = 1$	0.00	0.16	0.03	0.00	0.26	0.08	0.00	0.15	0.07	0.03	0.16	0.03	0.03
	$h = 12$	0.00	0.19	0.05	0.00	0.23	0.04	0.00	0.45	0.00	0.00	0.05	0.00	0.00

Table S.2: Relative frequency of occurrence of each model for OIL as best model, across vintages, according to: CRPS (rows 1-2), ACPS with  $c = 0.05$  (rows 3-4), ACPS with  $c = 0.95$  (rows 5-6)

		AR(1)	AR(1)-SV	AR(1)-tSV	AR(20)	AR(20)-SV	AR(20)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
CRPS	$h = 1$	0.03	0.04	0.06	0.07	0.04	0.06	0.10	0.04	0.06	0.05	0.11	0.17	0.18
	$h = 5$	0.06	0.05	0.08	0.11	0.09	0.11	0.09	0.07	0.07	0.07	0.05	0.07	0.08
ACPS 0.05	$h = 1$	0.03	0.11	0.05	0.04	0.09	0.05	0.12	0.05	0.06	0.05	0.13	0.11	0.10
	$h = 5$	0.03	0.08	0.07	0.10	0.19	0.12	0.14	0.09	0.04	0.04	0.05	0.03	0.02
ACPS 0.95	$h = 1$	0.04	0.12	0.05	0.05	0.10	0.05	0.11	0.03	0.06	0.04	0.12	0.12	0.10
	$h = 5$	0.05	0.13	0.05	0.16	0.19	0.09	0.12	0.07	0.03	0.02	0.04	0.02	0.02

Table S.3: Relative frequency of occurrence of each model for EEX as best model, across vintages, according to: CRPS (rows 1-2), ACPS with  $c = 0.05$  (rows 3-4), ACPS with  $c = 0.95$  (rows 5-6)

		AR(1)	AR(1)-SV	AR(1)-tSV	AR(7)	AR(7)-SV	AR(7)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
CRPS	$h = 1$	0.04	0.03	0.05	0.11	0.16	0.10	0.04	0.05	0.06	0.11	0.08	0.07	0.10
	$h = 7$	0.04	0.05	0.05	0.07	0.07	0.04	0.03	0.08	0.08	0.12	0.12	0.10	0.16
ACPS 0.05	$h = 1$	0.03	0.03	0.02	0.18	0.18	0.04	0.03	0.09	0.08	0.04	0.13	0.11	0.04
	$h = 7$	0.01	0.01	0.01	0.04	0.05	0.01	0.01	0.18	0.05	0.09	0.32	0.10	0.12
ACPS 0.95	$h = 1$	0.05	0.05	0.03	0.16	0.19	0.04	0.05	0.08	0.07	0.02	0.12	0.10	0.03
	$h = 7$	0.09	0.06	0.01	0.08	0.06	0.03	0.05	0.12	0.05	0.04	0.27	0.06	0.08

Table S.4: Ranking of probability forecasts and accuracy test. Best model, over vintages, according to: CRPS; ACPS with  $c = 0.05; 0.275; 0.5; 0.725; 0.95$  for the three different datasets: Employment (top); Oil (middle) and EEX (bottom).

EMPL													
Horizon 1	AR(1)	AR(1)-SV	AR(1)-tSV	AR(12)	AR(12)-SV	AR(12)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
ACPS(.,.;0.05)	12	9***	5***	13	2***	1***	11***	3***	10***	8***	4***	7***	6***
ACPS(.,.;0.275)	12	10***	9***	13	8***	7***	11***	2***	6***	5***	1***	3***	4***
ACPS(.,.;0.5)	13	10***	9***	11	8***	7***	12***	4***	6***	5***	1***	2***	3***
ACPS(.,.;0.725)	13	10***	9***	11**	2***	1***	12***	8***	6***	7***	4***	3***	5***
ACPS(.,.;0.95)	13	4***	3***	12	1***	2***	11***	5***	7***	9***	6***	8***	10***
CRPS	13	10***	9***	11	8***	7***	12***	4***	6***	5***	1***	2***	3***
Knuppel p-value	0.137	0.389	0.374	0.475	0.222	0.180	0.145	0.337	0.136	0.136	0.195	0.061	0.061
Horizon 12	AR(1)	AR(1)-SV	AR(1)-tSV	AR(12)	AR(12)-SV	AR(12)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
ACPS(.,.;0.05)	11	12	10	13	3	2	9***	4	6	5	1	8	7
ACPS(.,.;0.275)	12	8***	5***	13	1***	2***	11***	3***	7***	6***	4***	9***	10***
ACPS(.,.;0.5)	12	4***	3***	13	1***	2***	11***	5***	7***	8***	6***	9***	10***
ACPS(.,.;0.725)	12	4***	1***	13	2***	3***	11***	5***	6***	7***	8***	9***	10***
ACPS(.,.;0.95)	12	1***	2***	13	3***	4***	11***	5***	7***	8***	6***	9***	10***
CRPS	12	4***	3***	13	1***	2***	11***	5***	7***	8***	6***	9***	10***
Knuppel p-value	0.181	0.088	0.126	0.440	0.124	0.138	0.170	0.087	0.354	0.354	0.049	0.287	0.287
OIL													
Horizon 1	AR(1)	AR(1)-SV	AR(1)-tSV	AR(20)	AR(20)-SV	AR(20)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
ACPS(.,.;0.05)	12	8*	11	10**	7*	9**	13	6***	3***	4***	5***	2***	1***
ACPS(.,.;0.275)	12	8	13	7*	9	10	11	4***	5***	6***	1***	3***	2***
ACPS(.,.;0.5)	12	9	13	7	8	11	10	3***	6***	5***	1***	4***	2***
ACPS(.,.;0.725)	10	9	13	7	8	11	12	2***	6**	5**	1***	4***	3***
ACPS(.,.;0.95)	12	8	11**	9	7	10**	13	3**	6*	5*	1**	4*	2*
CRPS	12	9	13	7	8	11	10	3***	6***	5***	1***	4***	2***
Knuppel p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.011	0.042	0.042	0.001	0.159	0.159
Horizon 5	AR(1)	AR(1)-SV	AR(1)-tSV	AR(20)	AR(20)-SV	AR(20)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
ACPS(.,.;0.05)	10	5	7	9	3	6	12	13	2	8	11	1*	4*
ACPS(.,.;0.275)	5	8	7	6	4	3	11	1**	13	12	2*	10	9
ACPS(.,.;0.5)	3	7	6	8	4	5	11	1**	13	12	2	10	9
ACPS(.,.;0.725)	2	4	5	8	3	6	9	1	13	11	7	12	10
ACPS(.,.;0.95)	6	1	3*	8	2	4	13	5	10	7	11	12	9
CRPS	3	6	7	8	4	5	11	1**	13	12	2	10	9
Knuppel p-value	0.025	0.021	0.012	0.040	0.016	0.009	0.004	0.002	0.017	0.017	0.001	0.002	0.002
EEX													
Horizon 1	AR(1)	AR(1)-SV	AR(1)-tSV	AR(20)	AR(20)-SV	AR(20)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
ACPS(.,.;0.05)	11	8*	6***	9***	3***	1***	12	13	7**	4**	10	5**	2***
ACPS(.,.;0.275)	12	10	11	5***	2***	1***	13	9***	7***	6***	8***	4***	3***
ACPS(.,.;0.5)	12	10***	11*	7***	2***	5***	13	9***	6***	4***	8***	3***	1***
ACPS(.,.;0.725)	12	10***	11***	9***	5***	6***	13	8***	3***	2***	7***	1***	4***
ACPS(.,.;0.95)	12	7***	8***	11***	2***	6***	13	9***	3***	5***	10***	1***	4***
CRPS	13	10***	11*	7***	1***	5***	12	9***	6***	4***	8***	3***	2***
Knuppel p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Horizon 7	AR(1)	AR(1)-SV	AR(1)-tSV	AR(20)	AR(20)-SV	AR(20)-tSV	AR(1)-MS	TVP-AR(1)	TVP-AR(1)-SV	TVP-AR(1)-tSV	TVP-AR(2)	TVP-AR(2)-SV	TVP-AR(2)-tSV
ACPS(.,.;0.05)	5	13	12	1***	11	4	7	3	9	8	2	10	6
ACPS(.,.;0.275)	10	12	13	8***	9***	7***	11	6***	5***	2***	3***	4***	1***
ACPS(.,.;0.5)	10	12	13	9***	8***	7***	11	6***	4***	2***	5***	3***	1***
ACPS(.,.;0.725)	12	11**	10***	9***	8***	7***	13	5***	6***	2***	3***	4***	1***
ACPS(.,.;0.95)	12	11	10	9***	7*	5**	13	1***	8*	4*	2***	6*	3**
CRPS	10	12	13	9***	8***	7***	11	6***	4***	2***	5***	3***	1***
Knuppel p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Notes:  
<sup>1</sup> \*\*\*, \*\* and \* indicate scores are significantly different from 1 at 1%, 5% and 10%, according to the ACPStest in Section 2.2.  
<sup>2</sup> Gray cells indicate models that are correctly calibrated at 5% significance level according to the Knuppel test.

Centre for Applied Macroeconomics and Commodity Prices (CAMP)

will bring together economists working on applied macroeconomic issues, with special emphasis on petroleum economics.

BI Norwegian Business School  
Centre for Applied Macro - Petroleum economics (CAMP)  
N-0442 Oslo

[www.bi.no/camp](http://www.bi.no/camp)