

Proper Scoring Rules for Evaluating Density Forecasts with Asymmetric Loss Functions

Matteo Iacopini, Francesco Ravazzolo & Luca Rossini

To cite this article: Matteo Iacopini, Francesco Ravazzolo & Luca Rossini (2022): Proper Scoring Rules for Evaluating Density Forecasts with Asymmetric Loss Functions, Journal of Business & Economic Statistics, DOI: [10.1080/07350015.2022.2035229](https://doi.org/10.1080/07350015.2022.2035229)

To link to this article: <https://doi.org/10.1080/07350015.2022.2035229>



© 2022 The Author(s). Published with license by Taylor & Francis Group, LLC.



[View supplementary material](#)



Published online: 11 Mar 2022.



[Submit your article to this journal](#)



Article views: 959



[View related articles](#)



[View Crossmark data](#)

Proper Scoring Rules for Evaluating Density Forecasts with Asymmetric Loss Functions

Matteo Iacopini^{a,b}, Francesco Ravazzolo^{c,d,e}, and Luca Rossini^{f,g}

^aVrije Universiteit Amsterdam, Amsterdam, The Netherlands; ^bTinbergen Institute, Amsterdam, The Netherlands; ^cFree University of Bozen-Bolzano, Italy; ^dBI Norwegian Business School, Oslo, Norway; ^eRCEA, Rimini, Italy; ^fUniversity of Milan, Milan, Italy; ^gCa' Foscari University of Venice, Venice, Italy

ABSTRACT

This article proposes a novel asymmetric continuous probabilistic score (ACPS) for evaluating and comparing density forecasts. It generalizes 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. The (weighted) ACPS extends the symmetric (weighted) CRPS by allowing for asymmetries in the preferences underlying the scoring rule. A test is used 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 (U.S. employment growth) and of commodity prices (oil and electricity prices) with particular focus on the recent COVID-19 crisis period.

ARTICLE HISTORY

Received May 2021
Accepted January 2022

KEYWORDS

Asymmetric continuous probabilistic score; Asymmetric loss; Density forecast; Predictive distribution; Probabilistic forecast; Proper score; Weighted score

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 minimizing the expected squared forecast error (Gneiting 2011).

A universal approach to forecasting is the provision of a predictive density, known as probabilistic or density forecasting (see Elliott and Timmermann 2016, chap. 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). Probabilistic forecasts aim 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 functional thereof (e.g., mean or quantile). Several reasons have been suggested for preferring density over point forecasts (e.g., Elliott, Ghanem, and Krüger 2016). First, point forecasting is often associated with 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) for early works and the discussions in Timmermann (2006) and Gneiting (2011). Carriero, Clark, and Marcellino (2020)

extend the application of point forecasts to tail risk nowcasts of economic activity. Moreover, there is substantial interest in forecasting continuous variables outside economics, such as climate (Jasiński 2020), energy consumption (Adams and Shachmurove 2008), biomedical science and biology (Ioannidis 2009; Tripto et al. 2020). 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.

Asymmetry plays an important role in forecasting time series and in particular in examining the variation in the degree of asymmetry when the forecast horizon increases (see, e.g., Galbraith and van Norden 2019). However, a theoretical framework to test the asymmetry in density forecasting is missing and we contribute to this stream of literature by introducing a new asymmetric proper scoring rule, the ACPS.

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 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 a 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 concerning forecasts that give too high probability mass to the left part of the

CONTACT Matteo Iacopini  m.iacopini@vu.nl  Vrije Universiteit Amsterdam and Tinbergen Institute, Amsterdam, The Netherlands.

 Supplementary materials for this article are available online. Please go to www.tandfonline.com/UBES.

© 2022 The Authors. Published with license by Taylor & Francis Group, LLC.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

distribution (negative or low growth of employment). This is the case of the FED, which has recently fixed the target long-run unemployment rate around 4.1%.¹ With this objective, if an economic forecast points to a long-run unemployment rate higher than the 4.1% threshold, then the FED would probably intervene to lower it, whereas if the forecasted rate is below the threshold, it is likely that the FED may need to raise interest rates since the economy is overheating.

Other examples relate to energy markets that have recently experienced negative prices. West Texas Index (WTI) oil prices collapsed to -37.63 U.S. 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 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, to guide the process of making and assessing forecasts. Hence, we develop a measure that properly incorporates asymmetry in density forecasting evaluation, and we apply it to study forecasting asymmetry in the three above-mentioned datasets.

The main goal of this article is the proposal of novel and practical forecasting evaluation tools that can answer the increasing demand from policymakers and central bankers. We plan to achieve this result by introducing an innovative asymmetric scoring rule that can 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 emphasizing 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 article, 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* loss functions. In the latter, they illustrated 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, Timmermann, and Komunjer (2005), Elliott, Komunjer, and Timmermann (2008) 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 toward over-prediction for Japan, UK, and Italy, thus, the fiscal situation turns out to be better than expected. On the other hand, Canada's under-prediction takes place when the fiscal situation is worse than expected for Canada (negative forecast error) 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 Jannsen (2017) documented that the GDP growth forecasts made by professional forecasters tend to exhibit systematic errors, and tend to overestimate GDP growth. Moreover, Boero, Smith, and Wallis (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 judgments are based on an asymmetric loss. 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.

Concerning 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.

A natural way to evaluate and compare competing forecasts is the use of proper scoring rules, which assess calibration and sharpness simultaneously and encourage honest and careful forecasting. Specifically, a proper scoring rule is a function that compares a probabilistic forecast with a realization of the variable, such that it is maximized when the forecast corresponds to the true distribution generating the data. It is strictly proper if the maximum is unique. Despite the wide literature on the class of proper scoring rules for probabilistic forecasts of categorical and binary variables (see, e.g., 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 an 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 of reweighting the CRPS over the domain of the variable of interest, they inherit the symmetry of the latter. Diks, Panchenko, and Van Dijk (2011); Diks et al. (2014) propose an alternative method to compare the predictive accuracy of competing density forecasts on a specific region of interest, B (e.g., the tails of the density). The approach relies on a likelihood-based scoring rule that exploits the conditional likelihood (given that the actual observation lies in B) or the censored likelihood (with censoring of the observations outside B) and favors density forecasts that closely approximate the true density in the region of interest, B .

¹See <https://www.chicagofed.org/research/dual-mandate/dual-mandate>

Winkler (1994) did the first effort toward 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 article is 2-fold. 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 apply a Diebold–Mariano-type test to our ACPS 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 (U.S. 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 article 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 of 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 the evaluation of the performance of the forecasting density from both perspectives.

The rest of the article 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 and illustrates a comparison with the (weighted) CRPS in simulated experiments. Finally, Section 4 provides different applications for forecasting U.S. 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 summarizes 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 Ω . Let \mathcal{A} be a σ -algebra of subsets of Ω , and let \mathcal{P} be a convex class of probability measures on (Ω, \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 ω materializes, then the forecaster’s reward is $S(P, \omega)$.

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 $S(Q, Q) \geq S(P, Q)$, with equality holding if and only if $P = Q$.

The vast majority of the proper scoring rules proposed in the literature are symmetric (e.g., CRPS²), that is, their 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 toward 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 penalizes 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 $c < 0.5$ penalizes 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

²The continuous ranked probability score (Gneiting and Raftery 2007) is defined as

$$\text{CRPS}^*(P, y) = - \int_{-\infty}^{+\infty} (F(u) - \mathbb{I}(y \leq u))^2 du. \quad (1)$$

It is a proper scoring rule based on a symmetric (quadratic) loss function. In the following, we will use the negative orientation, that is $\text{CRPS} = -\text{CRPS}^*$.

probability score (ACPS) as

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

The following result shows the properness of our new score for every level of asymmetry. Our main contribution is a constructive proof that relies on the combination of some of the results in Matheson and Winkler (1976) and Winkler (1994) to obtain a scoring rule accounting for (i) continuous probability distributions, (ii) asymmetric loss, and (iii) being proper. Specifically, Matheson and Winkler (1976) are concerned with the definition of proper scoring rules for continuous probability distributions, whereas Winkler (1994) considers the problem of creating proper asymmetric scoring rules from symmetric ones, but it is limited to binary distributions. The constructive proof of Theorem 1 illustrates how to suitably combine the two approaches to get a scoring rule with the desired properties.

Theorem 1 (Properness). The asymmetric scoring rule ACPS defined in Equation (2) 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

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

where $c \in (0, 1)$ is the level of asymmetry. 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 (2). \square

The integrals in Equation (2) can be numerically approximated by truncating the domain to $[u_{\min}, y]$ and $[y, u_{\max}]$ such that

$$\begin{aligned} \text{ACPS}(P, y; c) &\approx \sum_{i=1}^N w_{2,i}^y (c^2 - P(u_{2,i}^y)^2) \left[\frac{1}{(1-c)^2} \mathbb{I}(P(u_{2,i}^y) > c) \right. \\ &\quad \left. + \frac{1}{c^2} \mathbb{I}(P(u_{2,i}^y) \leq c) \right] + \sum_{i=1}^N w_{1,i}^y ((1-c)^2 - (1-P(u_{1,i}^y))^2) \\ &\quad \left[\frac{1}{(1-c)^2} \mathbb{I}(P(u_{1,i}^y) > c) + \frac{1}{c^2} \mathbb{I}(P(u_{1,i}^y) \leq c) \right], \end{aligned} \quad (3)$$

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.³ 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,⁴ Example 1 reports the value of the score as a function of c , for several probabilistic forecasts. See the supplementary materials for further examples.

Example 1. Let us consider several Gaussian probabilistic forecasts P . In Figure 1 we show the value of the score on a range of asymmetry values $c \in \{0.05, 0.275, 0.50, 0.725, 0.95\}$, for a given observation y whose true density is a centered Gaussian with a standard deviation equal to 2. 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,

³ Let θ be the vector of all the model's parameters. The posterior predictive density is defined as

$$P(y_{t+1} | y_1, \dots, y_t) = \int_{\Theta} P(y_{t+1} | y_1, \dots, y_t, \theta) P(\theta | y_1, \dots, y_t) d\theta.$$

Unfortunately, the integral above cannot be analytically solved for many commonly used econometric models. However, the conditional and posterior distributions, $P(y_{t+1} | y_1, \dots, y_t, \theta)$ and $P(\theta | y_1, \dots, y_t)$, respectively, can usually be sampled from quite easily. Therefore, when adopting a Bayesian approach based on MCMC it is possible to circumvent the integration problem and obtain an approximation of the predictive distribution as follows. For each iteration $i = 1, \dots, M$ of the Gibbs sampler, one first gets a draw from the posterior distribution of the parameters, $\theta^{(i)} \sim P(\theta | y_1, \dots, y_t)$, then samples from the conditional distribution of the observation to get $y_{t+1}^{(i)} \sim P(y_{t+1} | y_1, \dots, y_t, \theta^{(i)})$. This results in a collection of M draws $\mathbf{y}_{t+1} = (y_{t+1}^{(1)}, \dots, y_{t+1}^{(M)})'$ from the posterior predictive distribution, allowing to evaluate any predictive feature of interest. See Koop (2003) for further details.

⁴ The parameter c is used to introduce asymmetry in the ACPS and can be thought of a proxy of the degree of asymmetry of the decision-maker's preferences. In those cases when the latter should be represented by functions whose asymmetry cannot be captured by a single parameter, the ACPS still provides a first-level approximation to them.

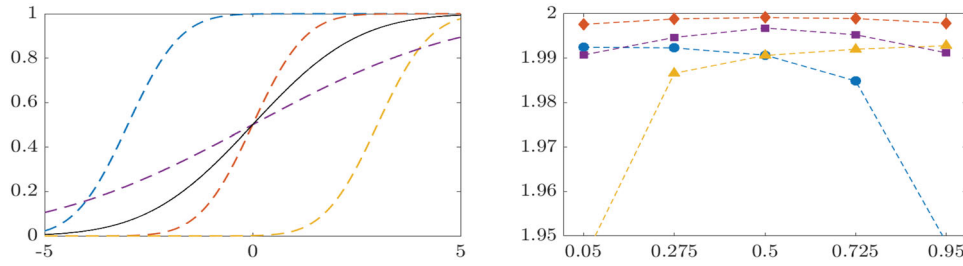


Figure 1. Asymmetric scoring rule $ACPS(P, y; c)$ for different forecasting densities P and asymmetry level c . The observed value is fixed at $y = 0$ and the true density is $\mathcal{N}(0, 4)$. Left panel: cumulative distribution functions of true density (solid, black) and forecasting densities: $\mathcal{N}(-3, 1)$ (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).

Table 1. Examples of weight functions for threshold-weighted and quantile-weighted CRPS, and variables supported on the real line.

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	$w(x) = 1 - \Phi(x)$	$v(\alpha) = (1 - \alpha)^2$

NOTE: ϕ, Φ denote the probability density and cumulative distribution functions of the standard Normal distribution, respectively, with $x \in \mathbb{R}$ and $\alpha \in (0, 1)$.

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 penalize 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.

2.1. Threshold and Quantile-Weighted Versions

In addition to asymmetric preferences toward 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 assigning 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, that is, asymmetric preferences, and domain-weighting of density forecasts.

Gneiting and Ranjan (2011) modified the CRPS by reweighting the loss according to a user-specified weight function, which allows selecting the regions where the decision-maker has a 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

$$tCRPS(P, y) = \int_{-\infty}^{+\infty} |P(z) - \mathbb{I}(y \leq z)|^2 w(z) dz, \tag{4}$$

$$qCRPS(P, y) = \int_0^1 2(\mathbb{I}(y \leq P^{-1}(\alpha)) - \alpha)(P^{-1}(\alpha) - y)v(\alpha)d\alpha, \tag{5}$$

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.

The definition of ACPS in (2) 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.⁵ We define the threshold-weighted asymmetric continuous probability score (tACPS), as

$$tACPS(P, y; c) = \int_{-\infty}^y (c^2 - P(u)^2) \left[\frac{1}{(1-c)^2} \mathbb{I}(P(u) > c) + \frac{1}{c^2} \mathbb{I}(P(u) \leq c) \right] G(du) + \int_y^{+\infty} ((1-c)^2 - (1-P(u))^2) \left[\frac{1}{(1-c)^2} \mathbb{I}(P(u) > c) + \frac{1}{c^2} \mathbb{I}(P(u) \leq c) \right] G(du), \tag{6}$$

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

$$qACPS(P, y; c) = \int_0^{P(y)} (c^2 - \alpha^2) \left[\frac{1}{(1-c)^2} \mathbb{I}(\alpha > c) + \frac{1}{c^2} \mathbb{I}(\alpha \leq c) \right] \frac{1}{p(P^{-1}(\alpha))} V(d\alpha) + \int_{P(y)}^1 ((1-c)^2 - (1-\alpha)^2) \left[\frac{1}{(1-c)^2} \mathbb{I}(\alpha > c) + \frac{1}{c^2} \mathbb{I}(\alpha \leq c) \right] \frac{1}{p(P^{-1}(\alpha))} V(d\alpha). \tag{7}$$

As stated for ACPS, we can provide evidence of the properness of the two novel scores defined in Equations (6) and (7).

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

1. The threshold-weighted asymmetric continuous probability score tACPS in Equation (6) is strictly proper;

⁵Notice that $G(du)$ is not required to be a probability measure.

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

		Domain	
		Uniform	Weighted
Loss	Symmetric	CRPS	tCRPS, qCRPS
	Asymmetric	ACPS	tACPS, qACPS

2. The quantile-weighted asymmetric continuous probability score $qACPS$ in Equation (7) is strictly proper.

Proof. From [Theorem 1](#), it is known that the ACPS is a proper scoring rule for any $c \in (0, 1)$, therefore, we are left to prove that a weighting scheme (threshold or quantile) preserves this feature. The result follows from the application of the procedure described in [Matheson and Winkler \(1976, Sec. 3\)](#), where the unweighed proper scoring rule is given by the ACPS. \square

Both tACPS and qACPS can be computed by approximating Equations (6) and (7) in a way analogous to Equation (3). The main advantage of the tACPS and qACPS consists in the ability 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).

[Table 2](#) provides a summary of some key differences between the CRPS and ACPS, and the corresponding weighted versions. We remark that the formula for the ACPS does not admit the CRPS as a special case. Instead, for $c = 0.5$, the (symmetric) ACPS has a similar interpretation to the CRPS when raking competing probabilistic forecasts. By the same token, the threshold- and quantile-weighted versions, tACPS and qACPS, for $c = 0.5$ can be interpreted similarly to the tCRPS and qCRPS measures.

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 (e.g., see [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. The density forecasting

approach requires a Diebold–Mariano-type test, since the forecast is an infinite dimensional object P .

Remark 2 (DM-type 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, then, the gain differential is defined as

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

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^* .

As claimed in [Diebold \(2015\)](#), when one is making model-based forecasts in settings where the true model is unknown, the DM test is approximately valid as long as its assumptions are approximately true. The DM test requires that the loss differential is covariance stationary, which means: for every t , $\mathbb{E}[d_t^*] = \mu$, $\text{cov}(d_t^*, d_{t-\tau}^*) = \gamma(\tau)$, and $\text{Var}[d_t^*] = \sigma^2 \in (0, \infty)$. The DM test has been extensively studied in the literature and some extensions have been proposed to improve its performance, for example in small samples, where parameter uncertainty does not vanish (see [Harvey, Leybourne, and Newbold 1997](#)).


Remark 3. Being a proper scoring rule, the ACPS can be used to compare and rank forecasts, in the spirit of the original Diebold–Mariano test ([Diebold 2015](#)). However, starting from [West \(1996\)](#); [Clark and McCracken \(2001\)](#), DM-type tests have been proposed for comparing models via forecasts, in pseudo-“out-of-sample” situations. This shift from forecast to model comparison requires to make assumptions not about the loss differential, but rather about the models, and ultimately results in the validity of the DM asymptotic standard normal null distribution depending on the nesting structure of the models ([West 2006](#); [Clark and McCracken 2013](#)). However, recent studies ([Clark and McCracken 2013](#)) have shown that standard normal critical values often approximate the exact null distribution very well, thus, supporting the use of these critical values in spite of alternative bootstrap procedures. Moreover, in the presence of model misspecification and parameter estimation error, the use of different scoring rules to rank competing models on the basis of the forecasting performance may result in different rankings ([Elliott, Ghanem, and Krüger 2016](#); [Patton 2020](#)).

In real-world applications, parameter estimation error may affect the forecasting results. To deal with this issue, in our forecasting exercise, we consider a rolling window approach where the length of the window used for estimation is substantially larger than that for out of sample comparison. We also remark that the empirical studies in this article are concerned with the investigation of the role of asymmetry in the decision maker’s preferences on the ranking of (possibly misspecified) models. Based on the PITs and calibration tests, we find evidence of



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

Figure 2. 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: $\mathcal{N}(0, 1)$ (black), $\mathcal{N}(-3, 1)$ (orange), $\mathcal{N}(3, 1)$ (yellow), $\mathcal{N}(0, 16)$ (purple).



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

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 $\mathcal{N}(2, 4)$ (black), forecasting densities are: $\mathcal{N}(0, 1)$ (blue), $\mathcal{N}(-3, 1)$ (orange), $\mathcal{N}(3, 1)$ (yellow), $\mathcal{N}(0, 16)$ (purple).

model misspecification, especially in the application to electricity prices (EEX dataset). Consistent with the previous literature, our findings on both synthetic and real-world data show that varying levels of the asymmetry parameter may yield different rankings of the competing models, in terms of their forecasting performance. Overall, as the ACPS is a proper scoring rule, this suggests that in presence of estimation errors and model misspecification the asymmetry of individual preferences guides the choice of the “best” model.

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 Gaussian target densities.⁶ 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.

Figure 2 provides graphical evidence of the properness of the ACPS in a Gaussian target. This figure 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.

To investigate further this aspect, Figure 3 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 (i.e., the one that maximizes the score), as does the ACPS for values of c around 0.5. However, the ranking significantly changes when the ACPS assigns more weight to

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

3.1. Threshold-Weighted Version

We deepen 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 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 Table 1, using 5 alternative asymmetry levels $c \in \{0.05, 0.275, 0.50, 0.725, 0.95\}$.

In Table 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 $c = 0.5$, that is, 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 favors 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 a heterogeneous weight to the performance on different intervals, while the latter asymmetrically rewards negative and positive deviations from the true value. The

⁶See the supplementary materials for different target densities example, such as Student-t, Gamma and Beta. This range includes families of distributions with different support (\mathbb{R}, \mathbb{R}_+ and $[0, 1]$), skewed and with fat tails.

Table 3. This table reports the ranking of probabilistic forecasts using tCRPS and tACPS, for different weights (uniform, center, tails, right and left tail) and asymmetry levels ($c \in \{0.05, 0.275, 0.5, 0.725, 0.95\}$).

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

NOTE: 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 $\mathcal{N}(0, 1)$, $\mathcal{N}(-3, 1)$, $\mathcal{N}(3, 1)$, $\mathcal{N}(0, 16)$.

threshold-weighted asymmetric scoring rule, tACPS, combines the two schemes and allows the evaluation of 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 toward under or overestimation, which motivates an asymmetric score.

4. Empirical Applications

In the empirical applications, we adopt a similar framework to Gneiting and Ranjan (2011), which noted that the weighted likelihood approach proposed in Amisano and Giacomini (2007) is not proper 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 in three applications related to macroeconomics (employment growth rate) and commodity prices (oil 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.

In this article, we have adopted the Bayesian paradigm for inference and relied on Markov chain Monte Carlo (MCMC) algorithms for the estimation of the parameters (see the supplementary materials for the details). Since the predictive densities of the competing models are not all available in closed form, we have followed the common practice in Bayesian statistics and

have obtained a sample from each predictive distribution along with the iterations of the Markov chain Monte Carlo algorithm (see also Remark 1).

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, 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.

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.⁷ 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.

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, Casarin, and Ravazzolo 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 their width is roughly correct on average, see Diebold, Gunther, and Tay (1998). The PITs indicate 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.

Table 4 shows the ranking of the probability forecasts over out-of-sample (OOS) windows and across models for all the three datasets for $c = 0.05, 0.5, 0.95$.⁸ The DM-type test of the

⁷Unfortunately, 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.

⁸See Table IV in the supplementary materials for results for a higher range of c .

Table 4. Ranking of probability forecasts and accuracy test.

		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***
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
Horizon 1	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
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***

NOTE: Best model, over OOS windows, according to: CRPS; ACPS with $c = 0.05; 0.5; 0.95$ for the three different datasets: Employment (top); Oil (middle) and EEX (bottom).
 1***, ** and * indicate scores are significantly different from the AR(1) benchmark at 1%, 5% and 10%, according to the DM-type test of the ACPs in Section 2.2.
 2 Bold numbers indicate models that are correctly calibrated at 5% significance level according to the Knuppel test.
 3 Gray cells indicate models that belong to the Superior Set of Models delivered by the MCS procedure at confidence level 10%.




Figure 4. Best model in each OOS window (computed over the previous 20 years of forecast) for EMPL dataset: CRPS (black), ACPS with $c = 0.05$ (red), ACPS with $c = 0.95$ (yellow).

ACPS presented in Section 2.2 is also reported.⁹ Moreover, we have employed the Model Confidence Set procedure of Hansen, Lunde, and Nason (2011) to jointly compare the predictive power of all models. We use the R package MCS detailed in Bernardi and Catania (2016) and differences are tested separately for each class of models (meaning for each panel in the tables and for each horizon) with a confidence level of $\alpha = 0.1$.

4.1. U.S. Employment Growth

In the first application, we aim at forecasting monthly U.S. total nonfarm seasonally adjusted employment growth rate downloaded from the FRED database. We consider the growth rate of the monthly employment rate in the United States from January 1980 to April 2020. We see evidence of some spikes, in particular with a strong fall in April 2020 due to the present COVID-19 situation (see Figure S.5 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.

The PIT tests in Table 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 the 12-month horizon, for which the p -value is marginally lower at 4.9%. Density forecasts from models TVP-AR(2)-SV¹⁰ and TVP-AR(2)-tSV are calibrated at 1-day ahead horizon; no density is correctly calibrated at 5-days ahead horizons.

Moreover, 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 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 modeling and forecasting (un)employment data. Moreover, the best model for $h = 12$ and $c = 0.5$ is the same when applying the CRPS. In Figure 4, we report the best model in each window for the two horizons ahead, where the black line refers to the CRPS, the red, and the yellow for the ACPS for $c = 0.05$ and $c = 0.95$, respectively. The graph shows large instability in the best model, in particular when using the CRPS. The ACPS rules

⁹In 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.

¹⁰Notice that the TVP-AR(2)-SV is always preferred in terms of relative accuracy.




Figure 5. Best model in each OOS window (computed over the previous 4 years of forecast) for OIL dataset: CRPS (black), ACPS with $c = 0.05$ (red), ACPS with $c = 0.95$ (yellow).

seem to prefer one of the alternative models for more consecutive OOS windows. 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 ages are found for other levels of c and h , despite model order varies substantially across measures.

4.2. West Texas Index

For oil prices, we analyze daily WTI data (no weekends) from January 02, 2012 to May 07, 2020 to include in the analysis the recent turmoil. 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 use a rolling window of 4 years and we forecast $h = 1$ and $h = 5$ days ahead using a recursive technique.

In the middle panel of Table 4, we find that across windows, for 1 day ahead the TVP-AR(2) is the best model whereas the TVP-AR(2)-SV is the second-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 a few models are superior to the benchmark. Figure 5 confirms that the ACPS is less variable in this selection than the CRPS.

Figure 6 illustrates the ACPS for one step ahead density forecasts of the OIL prices, according to a TVP-AR(2) model and an AR(20) model, for each OOS window of the rolling estimation and various levels of asymmetry. This figure presents some interesting insights. By looking at the scores between April 17 and April 21, we find that for both models 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 Figure 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, similar to the blue line in Figure 1.

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 both AR(20) and the TVP-AR(2) models in favor of alternatives for making forecasts. Conversely, an agent averse to underestimation facing the same decision problem, equipped with the same data and models, is likely to agree with one of the two models above.

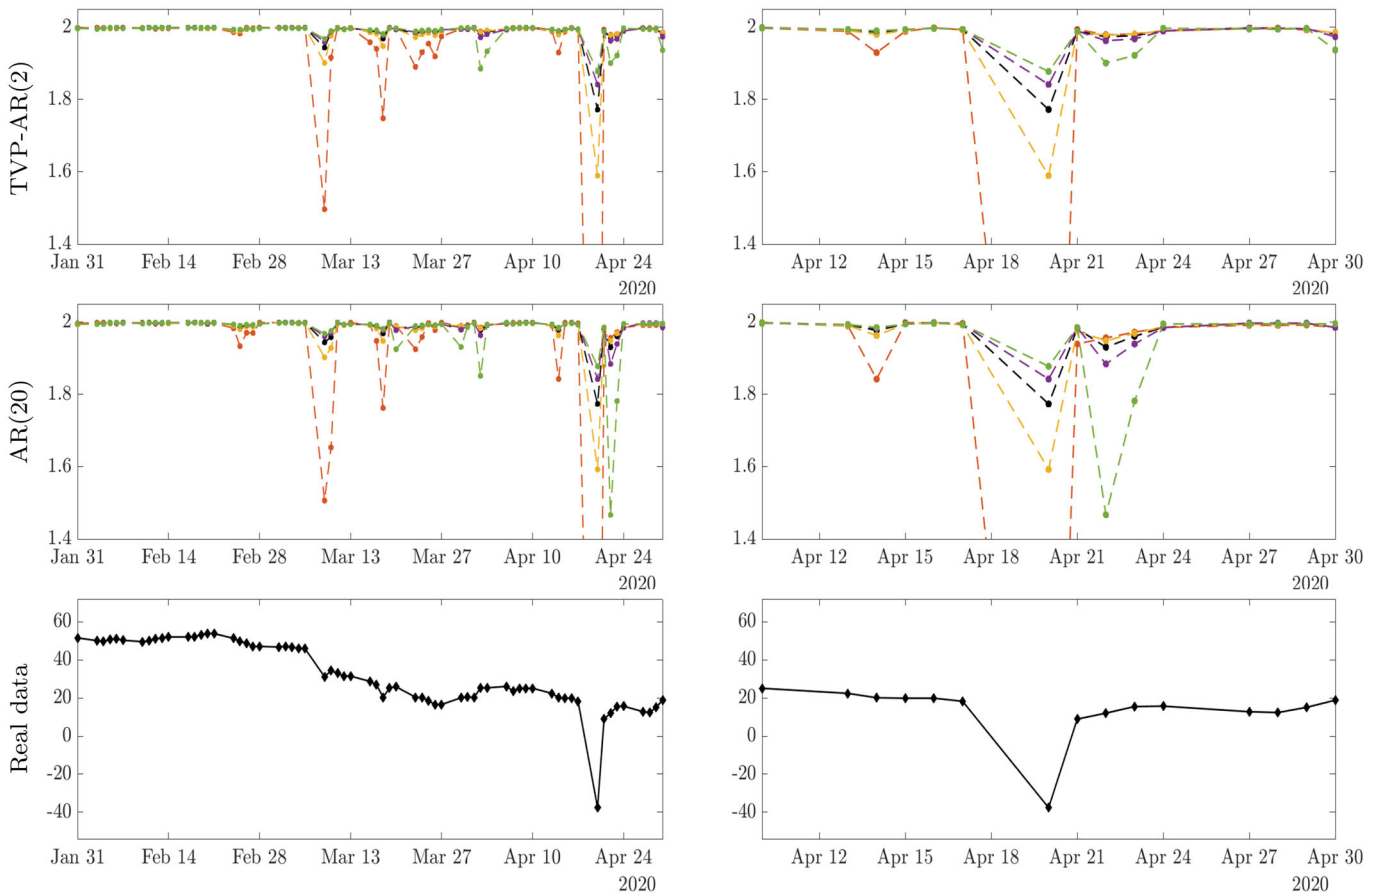


Figure 6. Top two rows: values of ACPS for one step ahead density forecasts of the OIL prices according to a TVP-AR(2) and an AR(20) model, respectively, for selected rolling windows (x-axis) and different asymmetry levels c : 0.05 (dashed red line), 0.275 (dashed yellow line), 0.50 (dashed black line), 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.

Moreover, these insights provide an important value-added of the ACPS as compared to symmetric scores. By looking at the 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.

4.3. Electricity Prices in Germany

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 materials). We analyze daily data (with weekends) from January 1, 2014 to May 8, 2020. For the forecasting analysis, we have considered a rolling window of 3 years and recursive techniques for predicting $h = 1$ and $h = 7$ days ahead.

From Table 4 we find that 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 ACPS 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 labeled as “purely econometric” models. Gianfreda, Ravazzolo, and Rossini (2020a) and Gianfreda, Ravazzolo, and Rossini (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.

The bottom panel of Table 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(7)-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

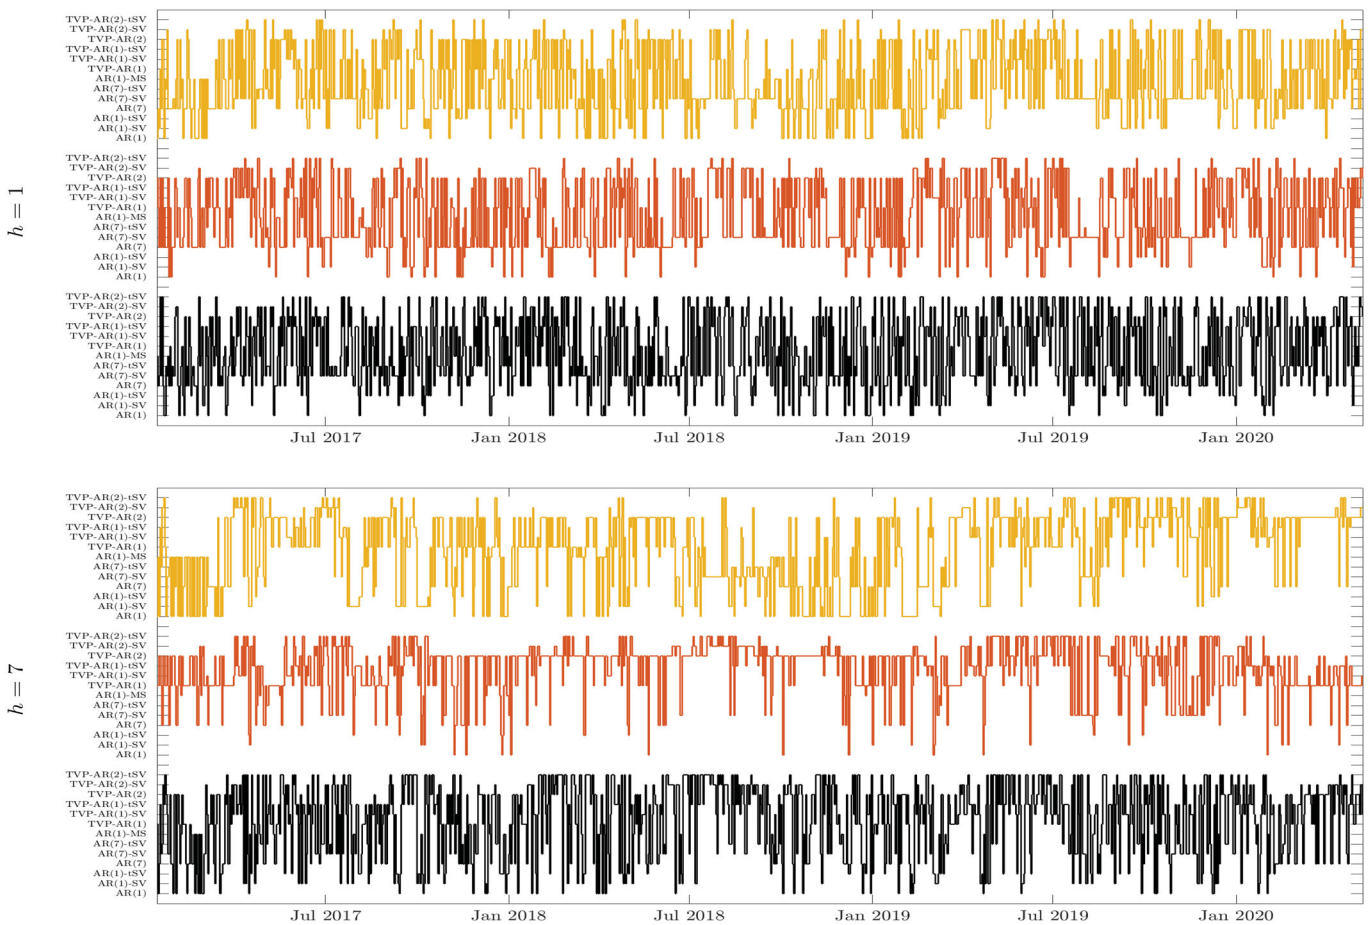


Figure 7. Best model in each OOS window (computed over the previous 3 years of forecast) for EEX dataset: CRPS (black), ACPS with $c = 0.05$ (red), ACPS with $c = 0.95$ (yellow).

in Gianfreda, Ravazzolo, and Rossini (2020b). At $h = 12$ the AR(7) 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. Figure 7 again indicates a more stable performance of some models when accounting for asymmetry relative to use the symmetric CRPS.

5. Conclusions

This article 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. Besides, 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 apply a DM-type 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 comparison.

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 a state-dependent function of economic variables, such as in Odendahl, Rossi, and Sekhposyan (2020).

Supplementary Materials

The supplementary material consists of four sections. Section S.1 illustrates different target densities for the simulation experiments. Section S.2 provides the details of the models used in the empirical applications. Finally, Section S.3 describes the data used in the analysis and Section S.4 provides additional details and results for the empirical analysis.

Acknowledgments

The authors are grateful to the Editor and the Reviewers for their useful comments which significantly improved the quality of the article. The authors gratefully acknowledge Giuseppe Cavaliere, Todd Clark, Michael McCracken, Massimiliano Marcellino, Barbara Rossi, Jonas Brehmer and seminar participants at University of Melbourne, Queen Mary University

of London, ESOBE 2021, IAEE 2021, 31st (EC)² Conference, 9th ICEEE 2021, 23rd Dynamic Econometrics Conference for their useful feedback. This article 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.

Funding

Matteo Iacopini acknowledges financial support from the EU Horizon 2020 programme under the Marie Skłodowska-Curie scheme (grant agreement no. 887220). Francesco Ravazzolo acknowledges financial support from Italian Ministry MIUR under the PRIN project “Hi-Di NET - Econometric Analysis of High Dimensional Models with Network Structures in Macroeconomics and Finance” (grant no. 2017TA7TYC).

References

- Adams, F. G., and Shachmurove, Y. (2008), “Modeling and Forecasting Energy Consumption in China: Implications for Chinese Energy Demand and Imports in 2020,” *Energy Economics*, 30, 1263–1278. [1]
- Amisano, G., and Giacomini, R. (2007), “Comparing Density Forecasts via Weighted Likelihood Ratio Tests,” *Journal of Business & Economic Statistics*, 25, 177–190. [8]
- 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. [1]
- Artis, M., and Marcellino, M. (2001), “Fiscal Forecasting: The Track Record of the IMF, OECD and EC,” *The Econometrics Journal*, 4, 20–36. [2]
- Bassetti, F., Casarin, R., and Ravazzolo, F. (2019), “Density Forecasting,” in *Macroeconomic Forecasting in the Era of Big Data*, eds. P. Fuleky, pp. 465–494, Cham: Springer. [8]
- Bauer, H. (2011), *Measure and Integration Theory* (Vol. 26), Berlin: Walter de Gruyter. [3]
- Bernardi, M., and Catania, L. (2016), “Portfolio Optimisation Under Flexible Dynamic Dependence Modelling,” arXiv e-prints. [10]
- Boero, G., Smith, J., and Wallis, K. F. (2008), “Evaluating a Three-Dimensional Panel of Point Forecasts: The Bank of England Survey of External Forecasters,” *International Journal of Forecasting*, 24, 354–367. [2]
- Carriero, A., Clark, T. E., and Marcellino, M. (2020), “Nowcasting Tail Risks to Economic Activity with Many Indicators,” Technical Report, Federal Reserve Bank of Cleveland, Working Paper No. 20-13. [1]
- Christodoulakis, G. A., and Mamatzakis, E. C. (2008), “An Assessment of the EU Growth Forecasts Under Asymmetric Preferences,” *Journal of Forecasting*, 27, 483–492. [2]
- (2009), “Assessing the Prudence of Economic Forecasts in the EU,” *Journal of Applied Econometrics*, 24, 583–606. [2]
- Christoffersen, P. F., and Diebold, F. X. (1996), “Further Results on Forecasting and Model Selection Under Asymmetric Loss,” *Journal of Applied Econometrics*, 11, 561–571. [2]
- Christoffersen, P. F., and Diebold, F. X. (1997), “Optimal Prediction Under Asymmetric Loss,” *Econometric Theory*, 13, 808–817. [2]
- Clark, T., and McCracken, M. (2013), “Advances in Forecast Evaluation,” in *Handbook of Economic Forecasting* (Vol. 2), eds. C. W. J. Granger, G. Elliott, and A. Timmermann, pp. 1107–1201, Amsterdam: Elsevier. [6]
- Clark, T. E., and McCracken, M. W. (2001), “Tests of Equal Forecast Accuracy and Encompassing for Nested Models,” *Journal of Econometrics*, 105, 85–110. [6]
- Dawid, A. P. (1982), “Intersubjective Statistical Models,” in *Exchangeability in Probability and Statistics*, eds. G. Koch and F. Spizzichino, pp. 217–232. Amsterdam: North-Holland Publishing Company. [8]
- Diebold, F. X. (2015), “Comparing Predictive Accuracy, Twenty Years Later: A Personal Perspective on the Use and Abuse of Diebold–Mariano Tests,” *Journal of Business & Economic Statistics*, 33, 1–9. [6]
- Diebold, F. X., Gunther, T. A., and Tay, A. S. (1998), “Evaluating Density Forecasts with Applications to Financial Risk Management,” *International Economic Review*, 39, 863–883. [8]
- Diks, C., Panchenko, V., Sokolinskiy, O., and van Dijk, D. (2014), “Comparing the Accuracy of Multivariate Density Forecasts in Selected Regions of the Copula Support,” *Journal of Economic Dynamics and Control*, 48, 79–94. [2]
- Diks, C., Panchenko, V., and Van Dijk, D. (2011), “Likelihood-Based Scoring Rules for Comparing Density Forecasts in Tails,” *Journal of Econometrics*, 163, 215–230. [2]
- Dovern, J., and Janssen, N. (2017), “Systematic Errors in Growth Expectations Over the Business Cycle,” *International Journal of Forecasting*, 33, 760–769. [2]
- Elliott, G., Ghanem, D., and Krüger, F. (2016), “Forecasting Conditional Probabilities of Binary Outcomes Under Misspecification,” *Review of Economics and Statistics*, 98, 742–755. [1,6]
- Elliott, G., Komunjer, I., and Timmermann, A. (2008), “Biases in Macroeconomic Forecasts: Irrationality or Asymmetric Loss?,” *Journal of the European Economic Association*, 6, 122–157. [2]
- Elliott, G., and Timmermann, A. (2004), “Optimal Forecast Combinations Under General Loss Functions and Forecast Error Distributions,” *Journal of Econometrics*, 122, 47–79. [2]
- (2016), *Economic Forecasting*. Princeton, NJ: Princeton University Press. [1]
- Elliott, G., Timmermann, A., and Komunjer, I. (2005), “Estimation and Testing of Forecast Rationality Under Flexible Loss,” *The Review of Economic Studies*, 72, 1107–1125. [2]
- Galbraith, J. W., and van Norden, S. (2019), “Asymmetry in Unemployment Rate Forecast Errors,” *International Journal of Forecasting*, 35, 1613–1626. [1]
- Giacomini, R., and Rossi, B. (2010), “Forecast Comparisons in Unstable Environments,” *Journal of Applied Econometrics*, 25, 595–620. [13]
- Gianfreda, A., Ravazzolo, F., and Rossini, L. (2020a), “Comparing the Forecasting Performances of Linear Models for Electricity Prices with High RES Penetration,” *International Journal of Forecasting*, 36, 974–986. [12]
- (2020b), “Large Time-Varying Volatility Models for Electricity Prices,” Technical Report, CAMP Working Paper Series 05/2020. [12,13]
- Gneiting, T. (2011), “Making and Evaluating Point Forecasts,” *Journal of the American Statistical Association*, 106, 746–762. [1]
- Gneiting, T., and Raftery, A. E. (2007), “Strictly Proper Scoring Rules, Prediction, and Estimation,” *Journal of the American Statistical Association*, 102, 359–378. [2,3]
- Gneiting, T., and Ranjan, R. (2011), “Comparing Density Forecasts Using Threshold- and Quantile-Weighted Scoring Rules,” *Journal of Business & Economic Statistics*, 29, 411–422. [2,5,8]
- (2013), “Combining Predictive Distributions,” *Electronic Journal of Statistics*, 7, 1747–1782. [1]
- Hansen, P. R., Lunde, A., and Nason, J. M. (2011), “The Model Confidence Set,” *Econometrica*, 79, 453–497. [10]
- Harvey, D., Leybourne, S., and Newbold, P. (1997), “Testing the Equality of Prediction Mean Squared Errors,” *International Journal of Forecasting*, 13, 281–291. [6]
- Ioannidis, J. P. (2009), “Limits to Forecasting in Personalized Medicine: An Overview,” *International Journal of Forecasting*, 25, 773–783. [1]
- Jasiński, T. (2020), “Use of New Variables Based on Air Temperature for Forecasting Day-ahead Spot Electricity Prices Using Deep Neural Networks: A New Approach,” *Energy*, 213, 118784. [1]
- Knoppel, M. (2015), “Evaluating the Calibration of Multi-Step-Ahead Density Forecasts Using Raw Moments,” *Journal of Business & Economic Statistics*, 33, 270–281. [8]
- Koop, G. (2003), *Bayesian Econometrics*, Chichester: Wiley. [4]
- Lerch, S., Thorarindottir, T., Ravazzolo, F., and Gneiting, T. (2017), “Forecaster’s Dilemma: Extreme Events and Forecast Evaluation,” *Statistical Science*, 32, 106–127. [5]
- Matheson, J. E., and Winkler, R. L. (1976), “Scoring Rules for Continuous Probability Distributions,” *Management Science*, 22, 1087–1096. [4,6]
- McCracken, M. W. (2020), “Diverging Tests of Equal Predictive Ability,” *Econometrica*, 88, 1753–1754. [6]
- Odendahl, F., Rossi, B., and Sekhposyan, T. (2020), “Comparing Forecast Performance with State Dependence,” Technical Report, Working Paper. [13]
- Patton, A. J. (2020), “Comparing Possibly Misspecified Forecasts,” *Journal of Business & Economic Statistics*, 38, 796–809. [6]

- Patton, A. J., and Timmermann, A. (2007), "Testing Forecast Optimality Under Unknown Loss," *Journal of the American Statistical Association*, 102, 1172–1184. [2]
- Rossi, B., and Sekhposyan, T. (2013), "Conditional Predictive Density Evaluation in the Presence of Instabilities," *Journal of Econometrics*, 177, 199–212. [8]
- Savage, L. J. (1971), "Elicitation of Personal Probabilities and Expectations," *Journal of the American Statistical Association*, 66, 783–801. [2]
- Schervish, M. J. (1989), "A General Method for Comparing Probability Assessors," *The Annals of Statistics*, 17, 1856–1879. [2]
- Timmermann, A. (2006), "Forecast Combinations," in *Handbook of Economic Forecasting* (Vol. 1), eds. C. W. J. Granger, G. Elliott, and A. Timmermann, pp. 135–196, Amsterdam: Elsevier. [1]
- Tripto, N. I., Kabir, M., Bayzid, M. S., and Rahman, A. (2020), "Evaluation of Classification and Forecasting Methods on Time Series Gene Expression Data," *Plos One*, 15, e0241686. [1]
- Tsuchiya, Y. (2016), "Assessing Macroeconomic Forecasts for Japan Under an Asymmetric Loss Function," *International Journal of Forecasting*, 32, 233–242. [2]
- West, K. D. (1996), "Asymptotic Inference About Predictive Ability," *Econometrica: Journal of the Econometric Society*, 64, 1067–1084. [6]
- (2006), "Forecast Evaluation," in *Handbook of Economic Forecasting* (Vol. 1), eds. C. W. J. Granger, G. Elliott, and A. Timmermann, pp. 99–134, Amsterdam: Elsevier. [6]
- Winkler, R. L. (1994), "Evaluating Probabilities: Asymmetric Scoring Rules," *Management Science*, 40, 1395–1405. [3,4]