

Discussion Paper

Deutsche Bundesbank
No 50/2022

Score-based calibration testing for multivariate forecast distributions

Malte Knüppel

(Deutsche Bundesbank)

Fabian Krüger

(Karlsruhe Institute of Technology)

Marc-Oliver Pohle

(Goethe University Frankfurt)

Editorial Board:

Daniel Foos
Stephan Jank
Thomas Kick
Martin Kliem
Malte Knüppel
Christoph Memmel
Hannah Paule-Paludkiewicz

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank,
Press and Public Relations Division, at the above address or via fax +49 69 9566-3077

Internet <http://www.bundesbank.de>

Reproduction permitted only if source is stated.

ISBN 978-3-95729-929-1

ISSN 2749-2958

Non-technical summary

Research Question

The word ‘forecast’ is often associated with the expected value of a certain variable. This type of forecast is a univariate point forecast. However, forecasts are increasingly issued in the form of probability distributions. From these, quantities like the expected value can be inferred. In addition, often these forecasts are simultaneously produced for several variables at once, for instance for economic activity, prices and interest rates. Such forecasts are called multivariate distributional forecasts. Any type of forecast should be investigated with respect to systematic errors. For example, univariate point forecasts are commonly checked for a systematic upward or downward bias. For multivariate distributional forecasts, up to now there are only a few, often very complicated approaches to find systematic errors.

Contribution

In contrast to the previous literature, we propose to use loss functions in order to detect systematic errors of multivariate distributional forecasts. Loss functions are usually employed only to compare the accuracy of different forecasts. The most popular loss function for univariate point forecasts is the squared forecast error. For multivariate distributional forecasts, there are similar loss functions, which are easy to apply. Our central idea is to employ these loss functions in order to compare the expected loss of a multivariate distributional forecast to its observed loss. Systematic differences between both quantities indicate the presence of systematic forecast errors.

Results

In simulations, our approach turns out to be well suited for detecting systematic errors of multivariate distributional forecasts. We demonstrate the usefulness of our approach in applications to macroeconomic and stock market forecasts. For both cases, simple and suitable possibilities to detect systematic forecast errors have not been available until now.

Nichttechnische Zusammenfassung

Fragestellung

In der Öffentlichkeit wird unter einer Prognose häufig der erwartete Wert einer bestimmten Variable verstanden. Bei einer solchen Prognose handelt es sich um eine univariate Punktprognose. Prognosen werden jedoch zunehmend in Form von Wahrscheinlichkeitsverteilungen erstellt, aus denen dann unter anderem der erwartete Wert hervorgeht. Zudem beziehen sich diese Prognosen oft auf mehrere Variablen gleichzeitig, also zum Beispiel auf wirtschaftliche Aktivität, Preise und Zinsen. Solche Prognosen werden multivariate Wahrscheinlichkeitsprognosen genannt. Jede Art von Prognosen sollte auf systematische Fehler überprüft werden. Bei univariaten Punktprognosen wird zum Beispiel oft untersucht, ob sie systematisch zu hoch oder zu niedrig ausfallen. Für multivariate Wahrscheinlichkeitsprognosen gibt es bisher nur wenige, oft sehr komplizierte Verfahren, um systematische Fehler zu finden.

Beitrag

Wir schlagen im Gegensatz zur bisherigen Literatur vor, Verlustfunktionen zu benutzen, um systematische Fehler von multivariaten Wahrscheinlichkeitsprognosen zu entdecken. Verlustfunktionen werden üblicherweise nur dazu verwendet, die Treffgenauigkeit verschiedener Prognosen miteinander zu vergleichen. Die bekannteste Verlustfunktion für univariate Punktprognosen ist der quadratische Prognosefehler. Auch für multivariate Wahrscheinlichkeitsprognosen gibt es einfach anzuwendende Verlustfunktionen. Unsere Kernidee besteht darin, diese Verlustfunktionen zu benutzen, um den erwarteten Verlust einer multivariaten Wahrscheinlichkeitsprognose mit dem tatsächlich beobachteten Verlust zu vergleichen. Systematische Abweichungen zwischen beiden Größen deuten auf systematische Prognosefehler hin.

Ergebnisse

Unsere Verfahren erweisen sich in Simulationen als gut geeignet, um systematische Fehler in multivariaten Wahrscheinlichkeitsprognosen zu entdecken. Wir zeigen die Nützlichkeit unserer Verfahren bei Anwendungen auf makroökonomische Prognosen und Finanzmarktprognosen. Für beide Fälle standen bisher keine einfachen, geeigneten Möglichkeiten zur Verfügung, um systematische Prognosefehler entdecken zu können.

Score-based calibration testing for multivariate forecast distributions*

Malte Knüppel[†]

Fabian Krüger[‡]

Marc-Oliver Pohle[§]

December 12, 2022

Abstract

Multivariate distributional forecasts have become widespread in recent years. To assess the quality of such forecasts, suitable evaluation methods are needed. In the univariate case, calibration tests based on the probability integral transform (PIT) are routinely used. However, multivariate extensions of PIT-based calibration tests face various challenges. We therefore introduce a general framework for calibration testing in the multivariate case and propose two new tests that arise from it. Both approaches use proper scoring rules and are simple to implement even in large dimensions. The first employs the PIT of the score. The second is based on comparing the expected performance of the forecast distribution (i.e., the expected score) to its actual performance based on realized observations (i.e., the realized score). The tests have good size and power properties in simulations and solve various problems of existing tests. We apply the new tests to forecast distributions for macroeconomic and financial time series data.

Keywords: Forecast Evaluation, Density Forecasts, Ensemble Forecasts

JEL classification: C12, C52, C53.

*We thank our reviewer Timo Dimitriadis as well as Tobias Fissler, Ana Galvão, Tilmann Gneiting, Johanna Ziegel and seminar participants at Deutsche Bundesbank, Heidelberg Institute of Theoretical Studies, Karlsruhe Institute of Technology, University of Cologne and participants of the International Symposium on Forecasting 2021 (virtual), the Statistische Woche 2021 (virtual), the Bundesbank/EABCN conference “Challenges in Empirical Macroeconomics since 2020” in Eltville, the DAGStat Conference 2022 in Hamburg, the International Association for Applied Econometrics 2022 Annual Conference in London, and the 2022 meeting of the Committee for Econometrics of the German Economic Association for helpful comments. We acknowledge support by the state of Baden-Württemberg through bwHPC. The views expressed in this paper are those of the authors and do not necessarily coincide with the views of the Deutsche Bundesbank or the Eurosystem.

[†]Deutsche Bundesbank, e-mail: malte.knuettel@bundesbank.de

[‡]Karlsruhe Institute of Technology, e-mail: fabian.krueger@kit.edu

[§]Goethe University Frankfurt, e-mail: pohle@econ.uni-frankfurt.de

1 Introduction

Decision-making and planning for the future requires good forecasts, often not only of the mean of a variable of interest, but also of other features of its probability distribution, for example the variance or certain quantiles. Ideally, probabilistic forecasts, that is forecasts of the full probability distribution, are provided. Consequently, there has been a gradual shift from point to probabilistic forecasting throughout scientific disciplines and fields of application (Gneiting, 2008; Gneiting and Katzfuss, 2014). At the same time, decision-making often requires forecasts of multiple variables of interest. Examples include forecasts of macroeconomic variables like GDP growth, inflation and an interest rate in economics, returns of a possibly large group of assets in finance or precipitation amounts at different locations in meteorology. Probabilistic forecasts then take the form of multivariate density or distribution functions. Such forecasts are, for instance, studied by Koop (2013), Chan (2020), McAlinn, Aastveit, Nakajima, and West (2020) and Carriero, Clark, Marcellino, and Mertens (2022) in macroeconomics, by Diks, Panchenko, Sokolinskiy, and van Dijk (2014), Catania, Grassi, and Ravazzolo (2019) and Gupta, Huber, and Piribauer (2020) in finance, and by Clark, Gangopadhyay, Hay, Rajagopalan, and Wilby (2004) and Heinrich, Hellton, Lenkoski, and Thorarinsdottir (2021) in meteorology.

When evaluating multivariate forecasts, it is crucial to use methods that evaluate the forecasts for all variables jointly. Evaluating each variable of interest separately would ignore the dependencies between them, which are usually of prime importance and the reason for multivariate forecasting in the first place. Forecast evaluation methods can be divided into methods for relative and absolute evaluation. Relative forecast evaluation is concerned with comparing the accuracy of different forecasts, while absolute forecast evaluation checks if a forecast fulfills certain desirable criteria. The comparison of multivariate probabilistic forecasts in terms of their accuracy, measured by the expected score, carries over straightforwardly from the univariate case. Popular proper scoring rules for multivariate forecast distributions are, for example, the log score and the energy score (Gneiting and Raftery, 2007). Absolute forecast evaluation usually amounts to examining calibration, that is the ‘statistical consistency between the distributional forecasts and the observations’ (Gneiting, Balabdaoui, and Raftery, 2007). For univariate probabilistic forecasts, calibration is commonly assessed by checking uniformity of probability integral transforms (PITs, see Dawid, 1984; Diebold, Gunther, and Tay, 1998; Gneiting et al., 2007). The PIT evaluates the forecaster’s cumulative distribution function (CDF) at the corresponding observation.¹ However, there is no multivariate analogue to the PIT uniformity result (Genest and Rivest, 2001). Two separate strands of literature (in econometrics on the one hand and meteorology and statistics on the other) try to generalize the PIT result from the univariate setting, but interestingly in quite different ways.

In this paper we propose simple and powerful tests of calibration for multivariate forecast distributions. We derive these tests from a general framework, which we introduce. This

¹For a formal expression of the PIT, see Definition 2.2 below.

framework enables us to construct tests that solve various problems of existing tests, and unifies the interdisciplinary literature in that it comprises many seemingly quite different tests. Before we provide an overview of our paper, we briefly review this literature.

1.1 Different strands of the literature

In econometrics, [Diebold et al. \(1998\)](#) and [Diebold, Hahn, and Tay \(1999\)](#) popularized an approach alluded to by [Smith \(1985\)](#), which uses the decomposition of the CDF of the d -dimensional forecast distribution $F(y) = F(y^{(1)}, y^{(2)}, \dots, y^{(d)})$ into conditional CDFs,

$$\begin{aligned}
 F(y^{(1)}, y^{(2)}, \dots, y^{(d)}) &= F(y^{(1)} | Y^{(2)} = y^{(2)}, \dots, Y^{(d)} = y^{(d)}) \\
 &\cdot F(y^{(2)} | Y^{(3)} = y^{(3)}, \dots, Y^{(d)} = y^{(d)}) \\
 &\cdot \dots \\
 &\cdot F(y^{(d-1)} | Y^{(d)} = y^{(d)}) \\
 &\cdot F(y^{(d)}),
 \end{aligned} \tag{1}$$

and the calculation of the PIT for each of the d conditional CDFs. Under the null hypothesis of optimality or ideal calibration, each PIT follows an independent $U(0, 1)$ distribution. While this insight is the starting point of virtually all multivariate calibration tests considered in the econometric literature, each of these tests has to cope with at least one of the following three problems.

First, assessing whether each PIT follows an independent $U(0, 1)$ distribution is not easy in practice. Jointly testing the distributional form and independence of each series might lead to power problems if d is large, and these types of tests have hardly been considered in the literature, with the approach of [González-Rivera and Yoldas \(2012\)](#) being one notable exception. [Diebold et al. \(1999\)](#) suggest to stack the d vectors containing the PITs and assess the distribution of the resulting vector. However, evaluating the independence property in a separate step gives rise to a multiple testing problem. Investigating the distribution of each marginal PIT, as sometimes done in the literature, aggravates this problem. Other approaches like those by [Clements and Smith \(2002\)](#), [Ko and Park \(2013\)](#) and [Dovern and Manner \(2020\)](#) avoid this problem by considering some function of the PITs and testing whether the resulting variable has the distribution implied under the null hypothesis.

Second, there is no natural ordering of the variables in the decomposition. The results of the tests described above hence depend on that ordering, which opens the floodgates to data snooping especially when d is large. [Dovern and Manner \(2020\)](#) provide details on this argument and propose order-invariant tests in response.

Third, tests based on (1) essentially require that the forecast distribution be available in closed form. When the forecast distribution is provided via a sample of draws, the tests are limited to small dimensions d , where reliable nonparametric estimation of the conditional CDFs is feasible.

The meteorological literature traditionally focuses on graphical diagnostic checks rather than on statistical testing. The usual diagnostic check for calibration of multivariate distributional forecasts proposed by [Gneiting, Stanberry, Gritmit, Held, and Johnson \(2008\)](#) consists of two steps: First, a so-called prerank function is applied to reduce the dimension from d to 1. Second, a univariate PIT or a rank histogram is constructed from these univariate quantities and checked visually for uniformity. Two popular prerank functions aim for a generalization of the univariate PIT: The multivariate rank histogram or Copula-PIT ([Gneiting et al., 2008](#); [Ziegel and Gneiting, 2014](#)) and the average rank histogram ([Thorarinsdottir, Scheuerer, and Heinz, 2016](#)). Like the tests in the econometric literature, they reduce to the univariate PIT for the case $d = 1$. [Wilks \(2017\)](#) and [Thorarinsdottir and Schuhen \(2018\)](#) provide overviews and comparisons of different prerank functions, while [Ziegel \(2015\)](#) discusses them from a theoretical perspective. [Gneiting et al. \(2008\)](#) also propose an alternative approach: They suggest to check standard uniformity of the Box density ordinate transform (BOT) by [Box \(1980\)](#).

Apparently, the only approach to testing calibration of multivariate probabilistic forecasts not making use of the PIT in some form stems from the biometric literature: [Wei, Balabdaoui, and Held \(2017\)](#) compute the expectation of the David-Sebastiani (1999) score for multivariate normal forecast distributions. In an unconditional test, they compare the average expected score under the forecast distribution to the average realized score. In an alternative, regression-based test they regress realized on expected scores. While theoretically appealing, the tests are only valid under restrictive assumptions under the null hypothesis, namely independent and identically distributed (IID) normal random vectors.

1.2 Contribution

We consider the null hypothesis of auto-calibration as introduced by [Tsyplakov \(2011\)](#) and [Gneiting and Ranjan \(2013\)](#) in the case of univariate probabilistic forecasts. Auto-calibration denotes optimality of the forecasts with respect to the information contained in the forecasts themselves. Intuitively, a forecast distribution is auto-calibrated if it cannot be improved by any form of transformation, for instance, by shifting its location or increasing its spread. The notion of auto-calibration generalizes readily to the multivariate case. This is in contrast to the weaker null hypothesis of probabilistic calibration ([Gneiting et al., 2007](#)), or equivalently, PIT uniformity, that is commonly employed in the univariate setting.

We introduce a general framework for testing auto-calibration of multivariate probabilistic forecasts. It is clear that information needs to be condensed in order to move from the null hypothesis involving multivariate forecast distributions and conditional distributions of outcomes to a one-dimensional test statistic. We make the choices involved in constructing the tests explicit and transparent. This sheds light on the relations between the different existing tests, most of which fall under this general framework, and uncovers which choices were implicitly made in the construction of those tests. Moreover, this setup facilitates the search for appropriate

choices and hence allows us to construct improved tests.

More specifically, we show that the construction of a test involves two steps: First, a dimensionality reduction function $g(F, Y)$ needs to be chosen. This function takes the outcome Y and the forecast distribution F as inputs and maps them to the real line, i.e. it reduces the dimensionality of the problem from d to 1. Note that we use the term ‘dimensionality reduction function’ rather than ‘prerank function’ because the former is more descriptive and is better adapted to our setup, which does not necessarily involve ranks. The second step then consists in the choice of a testing principle, with a PIT uniformity test being a natural choice: One tests whether the forecast distribution for the univariate quantity $g(F, Y)$ implied by F yields uniformly distributed PIT values. Many of the tests used in the literature (e.g., the tests by [Gneiting et al., 2008](#), [Thorarinsdottir et al., 2016](#), [Ziegel and Gneiting, 2014](#) and [Dovern and Manner, 2020](#)) fall under this PIT-based approach and merely differ in their choice of dimensionality reduction function. An alternative testing principle for the second step compares the realized value of the dimensionality reduction function, $g(F, Y)$, with the expected value under the forecast distribution, $\mathbb{E}_F[g(F, X)]$, where $X \sim F$, in terms of their means. This only requires a simple t -test. The regression-based test of [Wei et al. \(2017\)](#) is based on the same idea, but requires much more restrictive assumptions as mentioned above. Temporal dependence can be addressed by using heteroscedasticity-and-autocorrelation-consistent (HAC) approaches for both testing principles considered.

The choice of the dimensionality reduction function g is crucial for the performance of the tests. We propose to use proper scoring rules as dimensionality reduction functions because they are designed to summarize the quality of a forecast distribution and are thus sensitive to all of its facets. The resulting score-based calibration tests are powerful, straightforward to implement, interpretable and applicable in high-dimensional settings as we demonstrate in simulation studies and two case studies from macroeconomics (considering Bayesian vector autoregressive models) and finance (considering quantile regression and nonparametric copula methods).

1.3 Plan of the paper

Section 2 introduces our general framework for calibration testing for multivariate forecast distributions, while Section 3 describes the score-based tests. Section 4 presents simulation results, Section 5 and Section 6 contain case studies on macroeconomic and financial forecasting and Section 7 concludes. R-code to implement our proposed methods is available at <https://github.com/FK83/forecastcalibration>.

2 A general framework for testing calibration of multivariate forecast distributions

2.1 Auto-calibration and probabilistic calibration

Throughout the paper we denote by F_V the CDF of a random vector V , by $F_{V|\mathcal{H}}$ the (conditional) CDF of a random vector V conditional on the σ -algebra \mathcal{H} . By $\sigma(W)$ we denote the σ -algebra generated by a random variable W and when conditioning on $\sigma(W)$, we write $\bullet|W$ for $\bullet|\sigma(W)$. We abbreviate almost surely by *a.s.*

Consider a d -dimensional random vector of interest Y and the forecaster's information set \mathcal{I} , which is generated by a possibly very large vector of variables. A probabilistic forecaster's goal is to predict the conditional distribution of Y given \mathcal{I} , $F_{Y|\mathcal{I}}$. We simply write F for the forecast distribution itself.² For the purpose of forecast evaluation, a forecast-observation sample $\{(F_t, Y_t)\}_{t=1}^T$ is available.

We consider auto-calibration (Tsyplakov, 2011; Gneiting and Ranjan, 2013) as the null hypothesis of our tests.

Definition 2.1. *The forecast distribution F is **auto-calibrated** if it holds that*

$$F_{Y|F} = F \text{ a.s.}$$

A forecast is auto-calibrated if it is optimal relative to the information contained in itself, $\sigma(F)$. Intuitively, a user who receives an auto-calibrated forecast F should use this forecast 'as is', without any form of post-processing. Auto-calibration thus seems closely in line with the definition of calibration as the 'statistical consistency between the distributional forecasts and the observations' (Gneiting et al., 2007). The broader idea behind auto-calibration dates back to Theil (1961) and Mincer and Zarnowitz (1969) who consider a scatter plot of realized values Y (vertical axis) against mean forecasts \hat{Y} (horizontal axis), arguing that the observations in this diagram should scatter unsystematically around the 45 degree line. This requirement corresponds to $\mathbb{E}[Y|\hat{Y}] = \hat{Y}$, the analogue of Definition 2.1 for the mean functional, which can be viewed as the null hypothesis behind the widely-used Mincer-Zarnowitz (1969) regression. Tsyplakov (2011) and Gneiting and Ranjan (2013) introduce auto-calibration in the context of distributional forecasts and coin the term. Importantly, auto-calibration transfers to all types of forecast distributions F , whether univariate or multivariate, continuous or discrete.³

²Usually we are concerned with a time series forecasting problem where the forecaster stands at time $t - h$ and wants to forecast h periods into the future and \mathcal{I} represents all the information available at time $t - h$. Nevertheless, we do not use time indices in the theoretical part of this paper to avoid clutter and because the theory and tests presented here apply more generally to other classes of prediction problems as well, e.g. cross-sectional prediction problems. When we consider an evaluation sample, we will use the index t to refer to the forecast and observation for time t (or cross-sectional unit t), but still suppress the forecast horizon h .

³Furthermore, auto-calibration also transfers naturally to functionals such as quantiles and expectiles. Guler, Ng, and Xiao (2017) extend Mincer and Zarnowitz (1969) type regressions along these lines.

Probabilistic calibration introduced by [Gneiting et al. \(2007\)](#) reflects the current practice in the evaluation of univariate forecast distributions, in that checking PIT uniformity is the most common approach to absolute forecast evaluation throughout the disciplines. [Rosenblatt \(1952\)](#) presented the following definition and uniformity result of the PIT.

Definition 2.2. *Let H be a CDF and W a univariate random variable. Then*

$$PIT_{H,W} = H(W)$$

*is the **probability integral transform (PIT)** of H and W .*⁴

If $W \sim H$, the PIT uniformity result given by

$$PIT_{H,W} \sim U(0, 1)$$

holds. Note that the PIT readily generalizes to the multivariate case, but the uniformity result does not (see [Genest and Rivest, 2001](#)), i.e. the distribution of a random vector plugged into its own CDF is in general unknown.

Returning to the forecasting framework introduced above, in the univariate case the PIT uniformity result implies that under auto-calibration we have

$$PIT_{F,Y}|F \sim U(0, 1).$$

This conditional uniformity of the PIT implies unconditional uniformity, which has been called probabilistic calibration by [Gneiting et al. \(2007, Definition 1\(a\)\)](#) and [Gneiting and Ranjan \(2013, Definition 2.6\(b\)\)](#).

Definition 2.3. *The forecast distribution F is **probabilistically calibrated** if it holds that*

$$PIT_{F,Y} \sim U(0, 1).$$

Due to the lack of a PIT uniformity result in the multivariate case, probabilistic calibration does not generalize, but nevertheless most approaches in the literature on calibration testing for multivariate forecast distributions strive for a suitable extension of this concept. We take a different route by starting from the null hypothesis of auto-calibration and deriving testable implications. We thus rather view PIT uniformity as an implication of auto-calibration, as opposed to being a notion of calibration in its own right. This subtle shift of perspective suggests novel ways of constructing multivariate tests.

⁴The PIT can be adapted to non-continuous distributions by introducing randomization at the discontinuities (see e.g. [Rüschendorf \(2009\)](#)). We focus on the continuous case here, but note that all the results in this paper continue to hold in the general case. Note that in the general statistical literature, this generalization of the PIT to arbitrary distributions is called distributional transform (see e.g. [Rüschendorf, 2009](#)), but in the forecast evaluation literature it is usually still called PIT (see e.g. [Gneiting and Katzfuss, 2014, Definition 2](#)).

2.2 Testable implications of auto-calibration

The defining equation of auto-calibration from Definition 2.1 is hard to test directly as this would involve the estimation of a multivariate conditional distribution function conditional on a large information set. Thus, we are looking for testable implications, which requires to condense information. We get rid of the first difficulty by reducing the dimensionality from d to 1 by means of a dimensionality reduction function g which maps the forecast-observation pair (F, Y) to the real line. This step allows us to apply univariate evaluation methods to assess the forecast distribution of the univariate quantity $g(F, Y)$. The original forecast distribution F implies a forecast distribution for this auxiliary forecasting problem, which we call G . Let X be a random draw from F , i.e. $X \sim F$ (or equivalently $X = F^{-1}(U)$ with $U \sim U(0, 1)$). Then $G = F_{g(F, X)|F}$, that is G is a CDF-valued random variable fully determined by F . In the rest of the paper, often probabilities and expectations from G show up, where we are dealing with a specific g or X and want to keep g and X explicit. We then write compactly $\mathbb{P}_F(g(F, X) \leq v)$ instead of $\mathbb{P}(g(F, X) \leq v|F)$ for $G(v)$ and $\mathbb{E}_F[g(F, X)]$ for its expectation instead of $\mathbb{E}[g(F, X)|F]$.

Auto-calibration implies that

$$F_{g(F, Y)|F} = G \text{ a.s.} \quad (2)$$

As the estimation of $F_{g(F, Y)|F}$ is still practically infeasible due to the large information set $\sigma(F)$, we next get rid of the conditioning. There are two natural ways to achieve this. For the first, we use the PIT uniformity result and the law of total probability: Equation (2) can be written as

$$g(F, Y)|F \sim G \text{ a.s.},$$

which, by the PIT uniformity result, is equivalent to

$$PIT_{G, g(F, Y)}|F \sim U(0, 1) \text{ a.s.}$$

By the law of total probability, this implies that

$$U := PIT_{G, g(F, Y)} \sim U(0, 1) \text{ a.s.} \quad (3)$$

Conveniently, testing Equation (3) amounts to testing probabilistic calibration of the implied forecast distribution G . To this end, common tests of standard uniformity under temporal dependence can be applied.

Another implication of auto-calibration arises when we concentrate on the means of the distributions showing up in (2), which yields

$$\mathbb{E}[g(F, Y)|F] = \mathbb{E}_F[g(F, X)] \text{ a.s.} \quad (4)$$

Tsyplakov (2014) points out this implication and suggests to use it for testing in the case of univariate probabilistic forecasts. We now use the law of iterated expectations to get rid of the

conditioning:

$$\mathbb{E}[g(F, Y)] = \mathbb{E}[\mathbb{E}_F[g(F, X)]].$$

This equality represents an unbiasedness condition on the implied forecast distribution G . Rewriting it as

$$\mathbb{E}[D] = 0 \text{ with } D = g(F, Y) - \mathbb{E}_F[g(F, X)], \quad (5)$$

it is clear that it can be tested by a simple t -test using HAC standard errors.⁵

We call the two general approaches to testing auto-calibration of multivariate probabilistic forecasts coming out of equations (3) and (5) the *PIT-based approach* and the *t-test approach*. Most existing tests are special cases of these two approaches as we will discuss below.

2.3 Example


To illustrate the use of the quantities U and D from (3) and (5) for checking calibration, we consider a bivariate normal example with $Y \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \Sigma)$ and constant forecast distribution $F = \mathcal{N}(\mathbf{0}, 0.5 \times \Sigma)$, i.e. the actual variance is twice as large as the forecast variance. We further set $\Sigma = \begin{pmatrix} 1 & 0.5 \\ 0.5 & 1 \end{pmatrix}$. We use the log score as a dimensionality reduction function, $g(F, Y) = -\log(f(Y))$, where $f = F'$ and smaller scores correspond to more accurate forecasts. We obtain the distribution of g by simulating from the forecast distribution F . Details on these issues are discussed later. We draw a sample $\{Y_t\}_{t=1}^{10000}$. The left panel of Figure 1 shows a time series plot of (a simulation based estimate of) D_t as defined in (5).⁶ The random variable D_t represents the difference between the realized log score and the expected log score under the forecast distribution. In the present example, the expected score is constant (since the forecast distribution is the same in each period), whereas the realized score is a random variable. The figure shows that D_t has a positive mean (marked by the blue line), which is clear evidence against the zero mean condition from (5), that is against auto-calibration. More specifically, the positive mean indicates that the forecast performance in terms of the log score is worse than expected, which is in line with the fact that the forecast variance is too small. The right panel of Figure 1 visualizes the histogram of (a simulation based estimate of) U_t , i.e. of the PITs of the log score as defined in (3).⁷ The histogram displays a spike at large values $U_t > 0.9$. This is clear evidence against the uniformity condition from (3) and thus against auto-calibration. Furthermore, it indicates that the realized scores $S(F_t, Y_t)$ are systematically larger than the scores $S(F_t, X_t)$ simulated under the forecast distribution, again indicating a worse performance

⁵Alternatively, one could use tests directly arising from equation (4), i.e. run a linear regression of $g(F, Y)$ on a constant and $\mathbb{E}_F[g(F, X)]$ and test for the constant being equal to zero and the slope coefficient being equal to one. This would be a direct generalization of the approach of Wei et al. (2017) that does not rely on a specific distributional assumption or a specific score. This approach can be more powerful in certain situations as two parameters are estimated contrary to only one in a t -test arising from (5), but has the disadvantage that it leads to perfect collinearity if the entropy of the forecast distribution does not change and to multicollinearity if it hardly varies.

⁶Equation (13) describes the finite sample estimator of D_t .

⁷Equation (11) describes the finite sample estimator of U_t .

Time series plot of realized minus expected score



Histogram of score PIT




Figure 1: Illustration of the two test variants. Left: Time series plot of differences D_t between realized and expected log scores. Average difference indicated by blue line. Under auto-calibration, expected difference equals zero. Right: Histogram of PITs U_t of realized log scores. Each of the 20 bars has an expected height of 5 under auto-calibration.

than expected.

Of course, in a converse setting where the forecast variance is too large, we would observe a negative mean of D_t , as well as an accumulation of small values of U_t (i.e., realized scores would be systematically smaller than scores simulated under the forecast distribution).

2.4 Implementation of the general tests

Consider an evaluation sample $\{(F_t, Y_t)\}_{t=1}^T$ consisting of T forecast-observation pairs. The forecast distribution F_t might come in closed form as a CDF F_t itself or as the corresponding density f_t .⁸ The forecast distribution F_t might alternatively be represented by samples $\{X_{j,t}\}_{j=1}^J$, where each $X_{j,t} \in \mathbb{R}^d$ is distributed according to F_t . State-of-the-art forecasting methods increasingly take the latter form: For example, meteorological forecasts are often based on an ensemble of draws from a physics-based model (see [Vannitsem, Wilks, and Messner, 2018](#), for an overview). In economics, Bayesian prediction methods using MCMC sampling have become popular in the last decade (e.g. [Clark, 2011](#)). Simulation-based predictive distributions are particularly relevant in the multivariate case since flexible forms of dependence across variables may not easily be captured by parametric families. In the following, we allow for the case that the simulation samples are dependent across instances $j = 1, \dots, J$, as is common in MCMC samples, for example. However, we assume that the sequence of draws is strictly stationary and ergodic, with invariant distribution F_t . These assumptions are standard and ensure that the simulation output indeed represents the desired forecast distribution ([Krüger, Lerch, Thorarinsdottir, and Gneiting, 2021](#); [Craiu and Rosenthal, 2014](#)).

⁸Depending on the context we write \bar{F}_t when referring to the CDF in closed form, but also use F_t to denote the forecast distribution in general and not making explicit in which form it comes, e.g. when denoting a forecast-observation pair as (F_t, Y_t) .

If the forecast CDF F_t or PDF f_t is available in closed form, $g(F_t, Y_t)$ is typically straightforward to determine. If the calculation of g requires the forecast distribution in closed form, but the latter is given by a sample $\{X_{j,t}\}_{j=1}^J$, we can, of course, estimate F_t by the empirical CDF and f_t by a kernel density estimator. However, unless the dimension d is small, the curse of dimensionality will make this estimation unreliable, even if we can easily obtain large samples from the forecast distribution. In this case, it is hence preferable to use a dimensionality reduction function that does not require F_t or f_t in closed form. We provide examples below.

To calculate U_t from (3) or D_t from (5), the implied forecast distribution G_t of the dimensionality reduction function (evaluated at a specific point, namely $g(F_t, Y_t)$) or its expectation $\mathbb{E}_{F_t}[g(F_t, X_t)]$, respectively, are needed. G_t is usually not available in closed form, even if F_t is. This does not pose a problem as we usually can obtain a large number of draws from F_t , i.e. a sample $\{X_{j,t}\}_{j=1}^J$. This immediately provides us with a sample from G_t , $\{g(F_t, X_{j,t})\}_{j=1}^J$, from which we can calculate the empirical CDF, \hat{G}_t , which approximates G_t arbitrarily well with growing J if a Glivenko-Cantelli theorem holds. This is ensured by strict stationarity and ergodicity of $\{X_{j,t}\}$, which carries over to $\{g(F_t, X_{j,t})\}$. This in turn ensures that

$$\hat{U}_t = PIT_{\hat{G}_t, g(F_t, Y_t)}$$

approximates U_t arbitrarily well, i.e. $\hat{U}_t \xrightarrow{a.s.} U_t$ as $J \rightarrow \infty$. Similarly, a strong law of large numbers (also implied by strict stationarity and ergodicity) for $\{X_{j,t}\}$ ensures that

$$\hat{D}_t = g(F_t, Y_t) - \frac{1}{J} \sum_{j=1}^J g(F_t, X_{j,t}) \xrightarrow{a.s.} D_t. \quad (6)$$

To test the condition from (3), i.e. uniformity of U_t , one can use tests of standard uniformity under temporal dependence.⁹ Chen (2011) shows that many tests proposed in the literature can be regarded as moment-based tests. The class of moment-based tests like Knüppel's (2015) test are well-suited for this purpose, since HAC covariance matrix estimators can account for the potential presence of serial correlation in the underlying time series. These time series are transformations of U_t chosen such that their expectation equals the moment of interest, respectively. Therefore, moment-based tests typically only require that a central limit theorem holds for the sample means of these series. To test condition (5), we can use a t -test with HAC standard errors. Thus, our tests are straightforward to implement (e.g., by using the software accompanying this paper) and have correct asymptotic size under standard assumptions.

If forecasts are generated from an estimated model, the estimation error of the parameters can cause problems in the context of forecast evaluation, as shown by West (1996) for point forecasts. Chen (2011) explains how to address this problem for moment-based calibration tests of density forecasts. However, this approach is quite involved and hardly used. In virtually all applications, density forecasts are evaluated without taking the effects of parameter estimation

⁹Temporal dependence of U_t and D_t typically arises in multi-step forecasting.

error into account. This approach is rationalized by [Rossi and Sekhposyan \(2019\)](#) who argue that in this case, the forecast model and the estimation technique are evaluated jointly. See also [Giacomini and White \(2006\)](#) for a similar argument in the context of relative forecast evaluation.

3 Score-based calibration tests

3.1 Proper scoring rules as dimensionality reduction functions

Ideally, dimensionality reduction functions for calibration testing should fulfill four criteria. First, in order to yield a powerful test, $g(F, Y)$ should be sensitive to changes in F and Y . This criterion is not trivial as g condenses information from d dimensions into one dimension. A second criterion is applicability: A dimensionality reduction function may lead to powerful tests, but at the same time not be applicable in certain practically relevant situations. A main example here are dimensionality reduction functions that require the forecast distribution in closed form, which are not applicable when the forecasts are given as samples and the dimension is not small enough to make nonparametric estimation of the densities or CDFs feasible. A third criterion is interpretability: In case of a rejection of the null, do the tests, i.e. the distribution of $\{\hat{U}_t\}_{t=1}^T$ for the PIT-test or $\frac{1}{T} \sum_{t=1}^T \hat{D}$ for the t -test, provide some constructive information on what is wrong with the forecasts and how they can be improved? A fourth criterion is simplicity regarding the dimensionality reduction functions and the implementation of the associated tests. Practitioners may not be willing to use a test that is overly complicated to understand or implement. We argue that proper scoring rules fulfill these four criteria.

Proper scoring rules are the cornerstone of relative forecast evaluation for probabilistic forecasts, where they are used to rank competing forecasts in terms of expected scores. A proper scoring rule S maps a forecast-observation pair (F, Y) to the real line. Propriety is defined as follows ([Gneiting and Raftery, 2007](#)):

Definition 3.1. A *scoring rule* S is *proper* relative to a class of distributions \mathcal{P} if it satisfies

$$\int S(F_1, y) dF_1(y) \leq \int S(F_2, y) dF_1(y) \quad (7)$$

for all F_1, F_2 in \mathcal{P} . It is *strictly proper* if (7) holds with equality if and only if $F_1 = F_2$.

Hence if Y is distributed according to F_1 , stating F_1 yields a smaller (i.e., better) expected score than stating any other distribution F_2 . In that sense, propriety incentivizes honest and careful forecasting. Strictly proper scores are designed to take into account all facets of a multivariate forecast distribution and the corresponding observation: The optimal forecast is preferred over a forecast that is non-optimal with respect to any single facet. When comparing misspecified forecasts, proper scores should be sensitive to various facets. Thus, proper scoring rules can be expected to perform well with respect to power of the calibration tests, the first criterion mentioned in the previous section.

The most popular proper scoring rules for multivariate probabilistic forecasts are the log score,

$$\text{LS}(f, Y) = -\log(f(Y)),$$

and the energy score,

$$\text{ES}(F, Y) = \mathbb{E}_F\|X - Y\| - 0.5 \mathbb{E}_F\|X - X^*\|, \quad (8)$$

where $\|\cdot\|$ denotes Euclidean distance and X, X^* are two independent draws from F (Gneiting and Raftery, 2007; Székely and Rizzo, 2005). We use both scores in negative orientation, i.e. smaller scores indicate more accurate forecasts.

Theoretical considerations by Lerch, Thorarinsdottir, Ravazzolo, and Gneiting (2017), which relate Diebold-Mariano (1995) predictive accuracy tests using the log score to the Neyman-Pearson lemma (although in a simplistic setup), suggest that the log score leads to powerful tests when comparing forecast accuracy. We thus expect the log score to perform well in terms of power in our tests as well. Extensive simulation results by Ziel and Berk (2019) show that the energy score generally has good discrimination ability as compared to other multivariate scoring rules used in the literature as well. We will focus on these two specific scores for the rest of this paper due to their popularity and this evidence.

Regarding the second criterion, applicability, the log score requires the forecast density f in closed form. While using a kernel density estimate of the forecast density is possible in principle, this approach is challenging even in the univariate setup (Krüger et al., 2021) and thus seems unappealing in the multivariate case. The energy score on the contrary is widely applicable, in that its defining expectations in (8) can readily be estimated via empirical averages.

Proper scores as dimensionality reduction functions facilitate the interpretation of test results: In a nutshell, our testing approaches are based on comparing realized scores, $S(F, Y)$, to scores anticipated by the forecaster, $S(F, X)$ with $X \sim F$. If the former tend to be larger than the latter, the forecaster is overconfident in that she anticipates her forecasts to be better than they actually are.¹⁰ We detail the aspect of interpretability when discussing the two specific tests below. Finally, our tests based on proper scores are reasonably simple to understand and implement.

Besides their suitability in terms of sensitivity, applicability, interpretability and simplicity, there are further conceptual reasons to use proper scores as dimensionality reduction functions. As noted by Held, Rufibach, and Balabdaoui (2010) and Wei et al. (2017), the fact that scores are widely used in relative forecast evaluation suggests their usefulness for calibration testing as well. Furthermore, there is a connection between proper scores and calibration, in particular

¹⁰Of course, the evaluation of confidence is conditional on the score used. A simple example would be a normally distributed target variable and the issuance of a t -distributed forecast with correct mean and variance. These forecasts would turn out to be overconfident with respect to the log score, but neither over- or underconfident according to the Dawid-Sebastiani score, because the latter only evaluates the mean and variance forecast.

auto-calibration, via the Murphy decomposition of the expected score (Pohle, 2020; Gneiting and Resin, 2022; Dimitriadis, Gneiting, and Jordan, 2021). Finally, our use of scoring rules can be motivated from the perspective of statistical depth functions (Zuo and Serfling, 2000; Mosler and Mozharovskiy, 2022). The reason why the classical PIT uniformity result does not generalize to the multivariate case is that quantiles do not readily generalize to the multivariate case due to the lack of an order in \mathbb{R}^d (Genest and Rivest, 2001). The univariate PIT basically checks if all quantile levels are hit uniformly. Statistical depth functions, which create a center-outward ordering in \mathbb{R}^d , are often proposed as a multivariate generalization of quantiles. A generalization of probabilistic calibration along those lines would thus check if the observations hit all depths of the forecast distribution uniformly. Many depth functions have been proposed (see Mosler and Mozharovskiy, 2022). Thorarinsdottir et al. (2016) use the so-called band depth for graphical diagnostic checking of calibration. Using Euclidean depth, perhaps the most basic depth function, as the dimensionality reduction function $g(F, Y)$ in our PIT-based test is equivalent to using the energy score. The same holds true for the Mahalanobis depth and the Dawid-Sebastiani score.

3.2 Generalized Box transform test

When using the PIT-based testing approach with proper scoring rules as dimensionality reduction functions, the PIT of the implied forecast distribution from (3) becomes

$$U_S = \mathbb{P}_F (S(F, X) \leq S(F, Y)) \quad (9)$$

with $X \sim F$. For reasons explained in Section 3.4, we refer to tests based on U_S as *generalized Box transform (GBT)* tests.

Using the log score, (9) becomes

$$U_{LS} = \mathbb{P}_F (-\log (f(X)) \leq -\log (f(Y))). \quad (10)$$

As discussed in the previous section for a general g , we approximate the PIT of the implied forecast distribution U_{LS} by drawing a large sample, $\{X_{j,t}\}_{j=1}^J$, from F_t and calculating

$$\hat{U}_{LS,t} = \frac{1}{J} \sum_{j=1}^J \mathbf{1}(\log (f_t(X_{t,j})) \geq \log (f_t(Y_t))), \quad (11)$$

where $\mathbf{1}(\bullet)$ denotes the indicator function. Subsequently, we test standard uniformity of $\{\hat{U}_{LS,t}\}_{t=1}^T$ with an appropriate test. We use the test by Knüppel (2015) in our simulations and empirical case studies.

In contrast to the test based on the log score, the GBT test based on the energy score does not

require a closed form forecast distribution. Here, U_S from (9) becomes

$$U_{\text{ES}} = \mathbb{P}_F(\mathbb{E}_F\|X - X^*\| \leq \mathbb{E}_F\|X - Y\|),$$

where X, X^* are two independent draws from F .¹¹ In order to estimate U_{ES} , we use two distinct samples $\{X_{t,i}\}_{i=1}^{J_0}$ and $\{X_{t,j}^*\}_{j=1}^{J_1}$, both of which are drawn from F_t , and set

$$\hat{U}_{\text{ES},t} = \frac{1}{J_1} \sum_{j=1}^{J_1} \mathbf{1} \left\{ \frac{1}{J_0} \sum_{i=1}^{J_0} \|X_{t,i} - X_{t,j}^*\| \leq \frac{1}{J_0} \sum_{i=1}^{J_0} \|X_{t,i} - Y_t\| \right\}. \quad (12)$$

We then test standard uniformity of $\{\hat{U}_{\text{ES},t}\}_{t=1}^T$. In practice, it may sometimes be more convenient to use one sample and split it in the middle to obtain two samples than drawing two separate independent samples. This appears unproblematic if the dependence between the two samples is sufficiently weak.¹²

To interpret the test results or to check calibration more informally in a graphical manner, histograms of $\{\hat{U}_{\text{ES},t}\}_{t=1}^T$ can be used in a similar way as PIT histograms are used in the univariate case. If the height of the histogram bars increases over $[0, 1]$ as in the example in the previous section, the forecaster is overconfident as the realized scores tend to be higher than the ones expected by the forecaster. Vice versa, if the histogram decreases, the forecaster is underconfident. Further shapes of the histograms may also hint to certain deficiencies of the forecasts: For example a U-shape means that the forecaster usually tends to give an assessment of expected forecast performance in a medium range, while the true forecast performance often tends to be much better or much worse. Thus, while they may not be over- or underconfident, the forecasts have difficulties in distinguishing between situations of good and bad expected forecast performance (i.e., low and high expected scores) and are not that informative in that respect. Reversely, a hump-shaped histogram arises because the assessments of expected forecast performance tend to be more extreme than the true forecast performance.

3.3 Entropy test

Using proper scores as dimensionality reduction functions in our t -test approach based on (5) amounts to a comparison of the forecaster's expected score, $\mathbb{E}_F[S(F, X)]$, and the realized score, $S(F, Y)$, via their difference, $D = S(F, Y) - \mathbb{E}_F[S(F, X)]$. $S(F, Y)$ represents the 'ex

¹¹Using the energy score corresponds to the dimensionality reduction function $g(F, Y) = \mathbb{E}_F\|Y - X\|$. This function ignores the minuend in (8), which drops out when moving to the PIT-based or the t -test. This choice of g also establishes the connection between the use of the energy score and the Euclidean depth function mentioned above.

¹²Alternatively, the following estimator based on a single sample $\{X_{t,j}\}_{j=1}^J$ could be used:

$$\tilde{U}_{\text{ES},t} = \frac{1}{J} \sum_{j=1}^J \mathbf{1} \left(\frac{1}{J} \sum_{i=1}^J \|X_{t,i} - X_{t,j}\| \leq \frac{1}{J} \sum_{i=1}^J \|X_{t,i} - Y_t\| \right).$$

Simulation results suggest that this estimator performs similarly to the estimator in (12), but we use the latter throughout this paper.

post' performance of the forecast distribution F , with larger values corresponding to worse performance and higher prediction uncertainty of Y . The 'ex ante' expected score under F , $\mathbb{E}_F[S(F, X)]$, is known as the generalized entropy of F (see e.g. [Gneiting and Raftery, 2007](#)), which is why we call this testing approach the *entropy test*. Related comparisons between ex post (or 'objective') uncertainty and ex ante (or 'subjective') uncertainty have been considered for univariate forecast distributions by [Galvão and Mitchell \(2019\)](#), [Clements \(2014\)](#) and [Krüger and Pavlova \(2022\)](#) in the context of economic survey forecasts. The entropy test is easily interpretable: If $E[D] > 0$, realized uncertainty exceeds expected uncertainty, i.e. the forecaster is overconfident in an 'average' sense. The opposite holds if $E[D] < 0$. Further, plotting the difference in the two uncertainties, D , over time can yield additional insights beyond considering its mean only: This diagnostic plot may reveal periods of over- and underconfidence and thus point to deficiencies in a forecaster's ability to quantify uncertainty in certain situations.

For the log score, D from equation (5) becomes

$$D_{\text{LS}} = -\log(f(Y)) - \mathbb{E}_F[-\log(f(X))].$$

Thus, here the entropy test amounts to using a sample $\{X_{j,t}\}_{j=1}^J$ from F_t to compute

$$\hat{D}_{\text{LS},t} = -\log(f_t(Y_t)) + \frac{1}{J} \sum_{j=1}^J \log(f_t(X_{j,t})) \quad (13)$$

and then to perform a t -test with HAC standard errors for the null $E[D_t] = 0$ on $\{\hat{D}_{\text{LS},t}\}_{t=1}^T$.

For the energy score D is given by

$$D_{\text{ES}} = \mathbb{E}_F\|X - Y\| - \mathbb{E}_F\|X - X^*\|.$$

In line with Equation (12), this term can be estimated based on two distinct samples $\{X_{t,i}\}_{i=1}^{J_0}$ and $\{X_{t,j}^*\}_{j=1}^{J_1}$:

$$\hat{D}_{\text{ES},t} = \frac{1}{J_0} \sum_{i=1}^{J_0} \|X_{t,i} - Y_t\| - \frac{1}{J_0 J_1} \sum_{i=1}^{J_0} \sum_{j=1}^{J_1} \|X_{t,i} - X_{t,j}^*\|. \quad (14)$$

Alternatively, Equation 7 of [Gneiting et al. \(2008\)](#) suggests to approximate D_t by a single sample $\{X_{j,t}\}_{j=1}^J$ from F_t by setting

$$\tilde{D}_{\text{ES},t} = \frac{1}{J} \sum_{j=1}^J \|X_{t,j} - Y_t\| - \frac{1}{J^2} \sum_{i=1}^J \sum_{j=1}^J \|X_{t,i} - X_{t,j}\|. \quad (15)$$

The latter estimator is the analogue to the PIT estimator in Footnote 12.

3.4 Dimensionality reduction functions of existing approaches

Many of the existing tests are special cases of our PIT-based approach: They first apply a dimensionality reduction function and then test the condition from (3), i.e. standard uniformity of the PIT of the implied forecast distribution. We now discuss dimensionality reduction functions used in some of these tests, noting that there are more functions used especially in meteorology (e.g. Thorarinsdottir and Schuhen, 2018; Wilks, 2017).

Dovern and Manner (2020) propose two order-invariant tests based on the decomposition of F into conditional CDFs and the corresponding conditional PITs. Their preferred test uses the dimensionality reduction function

$$g_{DM}(F, Y) = \sum_{i=1}^d \left[\Phi^{-1} \left(PIT_{F_{Y^{(i)}|Y^{(1)}, \dots, Y^{(i-1)}, Y^{(i+1)}, \dots, Y^{(d)}, Y^{(i)}}} \right) \right]^2, \quad (16)$$

where Φ^{-1} is the quantile function of the standard normal distribution. The applicability of this test is limited to cases where the forecast distribution is given in closed form or where the dimension d is small enough so that a reliable nonparametric density estimation is feasible.¹³

The dimensionality reduction functions showing up in the meteorological literature, where they are rather used for graphical diagnostic checking than for hypothesis testing, do not use conditional PITs. Nevertheless, most of them employ the PIT in some form. The multivariate rank histogram or Copula PIT (Gneiting et al., 2008; Ziegel and Gneiting, 2014) uses the multivariate PIT as dimensionality reduction function:

$$g_{CoPIT}(F, Y) = F(Y).$$

However, as the \leq -relation on \mathbb{R}^d only provides a partial order, the event $\{X \leq Y\}$ is very uninformative even for moderate dimensions d . Hence $F(Y) = \mathbb{P}_F(X \leq Y)$ will often be close to zero, making it hard to distinguish between an auto-calibrated forecast distribution F and miscalibrated alternatives.¹⁴ This leads to tests based on the Copula PIT having low power, a problem getting worse with increasing d .

A simple approach proposed by Thorarinsdottir et al. (2016) is the average rank histogram or average PIT, which just averages over all d marginal PITs:

$$g_{AvPIT}(F, Y) = \frac{1}{d} \sum_{i=1}^d F^{(i)}(Y^{(i)}). \quad (17)$$

This approach does not explicitly employ information contained in F beyond its marginals. We consider it in our simulations in the next section and shed some light on its usefulness there.

¹³Under the assumption that Y is multivariate normal, G , i.e. the distribution of $g_{DM}(F, X)$ under the null hypothesis, is known and U can be calculated directly without using draws from the forecast distribution.

¹⁴Note that here we compactly write $\mathbb{P}_F(X \leq Y)$ for $\mathbb{P}(X \leq Y|F)$ as well.

Remarkably, the Box transform as introduced by [Gneiting et al. \(2008\)](#) for checking calibration graphically does not use the PIT in the dimensionality reduction function $g(F, Y)$. It is defined as

$$U_{BT} = 1 - \mathbb{P}_F(f(X) \leq f(Y)), \quad (18)$$

where f is the forecast density corresponding to F , and $X \sim F$. Note that U_{BT} is equivalent to U_{LS} from (10), since the logarithm is a monotonic function. Thus, our PIT-based test with a proper score as a dimensionality reduction function can be viewed as a generalization of testing with the Box transform, hence the name ‘generalized Box transform (GBT) test’.

In our simulations, we consider the Dovern-Manner test and a test based on the average PIT. Concerning the four criteria mentioned above, the Dovern-Manner test will turn out to be appealing in terms of power, but its applicability is limited to the case of closed-form or low-dimensional forecast distributions. Furthermore, the test is hard to interpret (as the dimensionality reduction function from (16) averages over many terms involving conditional PITs) and to implement. In contrast to the Dovern-Manner test, tests based on the average PIT have very limited power in our simulations, but are relatively simple to interpret and easy to implement. Using the copula PIT as a dimensionality reduction function seems difficult to motivate in terms of either criterion.

4 Simulation studies

We next compare various tests via simulation, closely following the design of [Dovern and Manner \(2020, Section 3.1\)](#). The null hypothesis H_0 states that the data follows a multivariate normal distribution where each of the d variables has a mean equal to 0, a variance equal to 1, and its correlation with each other variable equals 0.5. Simulating data from this distribution for different dimensions d and different numbers of periods T will inform us about the size of the tests considered. To evaluate the tests’ power, we use the multivariate forecast distribution of H_0 , and data generated by four alternative distributions:

- H_1 : Multivariate normal distribution, variance of each variable equals 1.1^2 , all other parameters (i.e., means and correlations) like under H_0 .
- H_2 : Multivariate normal distribution, correlation of each pair of variables equals 0.4, all other parameters (i.e., means and variances) like under H_0 .
- H_3 : Multivariate t distribution with 8 degrees of freedom, rescaled such that covariance is like under H_0 .
- H_4 : Multivariate Gaussian constant conditional correlation (CCC)-GARCH(1,1) model of [Bollerslev \(1990\)](#) with GARCH parameters $\omega = 0.05, \alpha = 0.1, \beta = 0.85$; means and unconditional covariance matrix like under H_0 .

Note that our setup only differs from [Dovern and Manner \(2020\)](#) with respect to H_3 , where we rescale the distribution such that differences from H_0 only occur in fourth (and larger even) moments. This makes it easier to distinguish the results from those under H_1 , where ‘under H_i ’ means that the forecaster uses the distribution of H_0 , while the data-generating process (DGP) is given by H_i .

We only consider the $Z_t^{2\dagger}$ test of [Dovern and Manner \(2020\)](#), since applying their Z_t^{2*} test would be too time-consuming even in the case of normal forecast densities unless d is small. Moreover, both tests have similar size and power properties. We also employ the test based on the average-rank histogram proposed by [Thorarinsdottir et al. \(2016\)](#). For the entropy tests, the difference \hat{D}_t between the realized score and the expected score is regressed on a constant, and the null hypothesis of this constant being equal to zero is tested. $\hat{D}_t = \hat{D}_{ES,t}$ is described in equation (14) for the case of the energy score. For the case of the log score, $\hat{D}_t = \hat{D}_{LS,t}$ is given in (13). For the tests based on the generalized Box transform (GBT), we use Knüppel’s test.¹⁵ The PITs entering the test are given by (12) for the energy score and by (11) for the log score. As implied by (17), the average rank PITs are calculated as

$$\hat{U}_{Av,t} = \frac{1}{J} \sum_{j=1}^J \mathbf{1} \left(\frac{1}{d} \sum_{i=1}^d F \left(X_{t,j}^{(i)} \right) < \frac{1}{d} \sum_{i=1}^d F \left(Y_t^{(i)} \right) \right).$$

The results in Table 1 show that all tests considered have good size properties. With respect to power, however, pronounced differences arise. The Dovern-Manner test and the log-score-based GBT test tend to attain relatively high power in all situations considered, and their power increases with the number of variables d and the number of periods T . Power also increases with d and T for the energy-score-based GBT test. The power of the average rank test increases with T , but often remains fairly constant or even decreases when d increases. This behavior can occur if the individual ranks have the same correlation under the null and the alternative hypothesis, but different variances. These conditions turn out to hold (at least approximately) under H_1 , H_3 and H_4 .¹⁶ The entropy tests can also suffer from a lack of power, but at least in the case of the log score under H_3 , this result has a simple explanation. The expected score in this case only depends on the mean vector and the covariance matrix of the forecast

¹⁵[Dovern and Manner \(2020\)](#) employ [Neyman’s \(1937\)](#) smooth test, but this test is hardly applied in the corresponding literature. In the simulation design of [Dovern and Manner \(2020\)](#), [Neyman’s \(1937\)](#) smooth test has more power than Knüppel’s test, but the opposite would hold if, for instance, the true variances would be smaller than 1 under H_1 .

¹⁶In all situations considered, the mean of the average rank equals 0.5. The variance of the average rank is given by $\sigma^2(\rho + \frac{1-\rho}{d})$, where σ^2 denotes the variance of an individual rank and ρ denotes the correlation between two distinct individual ranks. This simple formula holds because in each DGP used, σ^2 and ρ are constant across all individual ranks and pairs of distinct individual ranks, respectively. Letting $\sigma_{H_0}^2$ and ρ_{H_0} denote these quantities under H_0 , and letting $\sigma_{H_1}^2$ and ρ_{H_1} denote these quantities under H_1 , two interesting special cases occur. If $\rho_{H_1} = \rho_{H_0} = \bar{\rho}$, the variances of the average ranks differ by $(\bar{\rho} + \frac{1-\bar{\rho}}{d})(\sigma_{H_1}^2 - \sigma_{H_0}^2)$. Obviously, this difference *decreases* (in absolute terms) with the number of variables d , as observed in the cases H_1 and H_4 . On the other hand, if $\sigma_{H_1}^2 = \sigma_{H_0}^2 = \bar{\sigma}^2$, the difference becomes $\bar{\sigma}^2(1 - \frac{1}{d})(\rho_{H_1} - \rho_{H_0})$, which *increases* (in absolute terms) with the number of variables d , as observed in the case H_2 . While these mechanisms are important to understand the behavior of the average rank test, they cannot explain every aspect, since higher moments like kurtosis also play a role.

distribution. Since the true distribution has the same mean vector and covariance matrix, the expected score and the average realized score are identical.¹⁷ Actually, using the log score and a normal forecast distribution under H_0 and H_3 , conducting the entropy test amounts to measuring the difference of the respective averages of the Mahalanobis distance from zero in this setup. The energy score is based on the Euclidean distance instead. The low power of the energy-score-based entropy test under H_3 thus indicates that the expected Euclidean distance under H_0 and the average realized Euclidean distance are similar. The same holds under H_1 . Under H_2 , due to the larger variance of the DGP, the expected distance is obviously smaller than the average realized distance, be it of the Mahalanobis or the Euclidean type. Accordingly, both entropy tests have relatively high power in this case.

Summing up, in the simulation setup based on [Dovern and Manner \(2020\)](#), the GBT tests turn out to be powerful tools to assess the calibration of multivariate forecast densities. The entropy tests do not have power against certain alternatives, but their results are easier to interpret, as done above for their behavior under H_1 , H_2 and H_3 . Finally, it is worth stressing that all tests except for the Dovern-Manner test are fairly easy to apply to non-Gaussian distributions. If the forecast density is unknown or hard to approximate, but a sample from this distribution is available, only the average rank test and the energy-score-based tests can be used.

Additional simulation results for the GBT and entropy test covering time series dependence, t -distributed forecast distributions and heteroskedasticity can be found in [Appendix A](#).

5 Macroeconomic case study: Forecasting US growth, inflation, and interest rates

We next apply our calibration tests to forecasts of three quarterly US macroeconomic variables: The gross domestic product (GDP) growth rate, the inflation rate as measured by the GDP deflator, and the three-month interest rate (T-Bill rate).¹⁸ We consider forecast horizons of $h = 0, 1, \dots, 4$ quarters ahead, and our evaluation sample ranges from 1965:Q4 to 2021:Q3 (target dates).

We employ a time-varying vector autoregressive model as proposed by [Primiceri \(2005\)](#) and [Del Negro and Primiceri \(2015\)](#) to generate forecasts, which is in line with [Dovern and Manner \(2020\)](#).¹⁹ The model is estimated via Bayesian methods, where we follow the default implemen-

¹⁷Note that the same result would be obtained using the score proposed by [Dawid and Sebastiani \(1999\)](#), which considers the first two moments only.

¹⁸For GDP and the GDP deflator, we use real-time data as provided by the Philadelphia Fed (series codes ROUTPUTQvQd and PQvQd). We consider logarithmic growth rates, and define the outcome as the second available data vintage. For the interest rate, we consider un-revised data from the FRED data base (series code DTB3).

¹⁹The corresponding application of the their tests requires challenging nonparametric kernel estimation procedures.

d	$T = 50$						$T = 200$					
	$Z_t^{2\ddagger}$	AvR	ES _D	ES _{GBT}	LS _D	LS _{GBT}	$Z_t^{2\ddagger}$	AvR	ES _D	ES _{GBT}	LS _D	LS _{GBT}
	Size											
2	0.05	0.06	0.07	0.05	0.07	0.05	0.06	0.05	0.05	0.05	0.05	0.05
4	0.05	0.05	0.06	0.05	0.05	0.05	0.06	0.06	0.05	0.05	0.05	0.05
6	0.05	0.05	0.06	0.05	0.06	0.05	0.05	0.05	0.05	0.04	0.04	0.05
10	0.06	0.05	0.06	0.05	0.05	0.06	0.06	0.05	0.06	0.05	0.06	0.05
20	0.05	0.05	0.06	0.06	0.05	0.05	0.05	0.05	0.06	0.05	0.05	0.05
50	0.05	0.05	0.06	0.06	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05
	Power given data from H_1 (larger variance)											
2	0.12	0.05	0.16	0.11	0.17	0.13	0.44	0.14	0.65	0.45	0.71	0.51
4	0.21	0.05	0.25	0.17	0.37	0.24	0.79	0.12	0.85	0.69	0.95	0.86
6	0.31	0.05	0.31	0.22	0.54	0.36	0.94	0.11	0.92	0.81	0.99	0.97
10	0.55	0.05	0.37	0.30	0.80	0.59	1.00	0.11	0.97	0.93	1.00	1.00
20	0.90	0.05	0.44	0.45	0.98	0.91	1.00	0.10	0.99	0.99	1.00	1.00
50	1.00	0.05	0.51	0.76	1.00	1.00	1.00	0.10	1.00	1.00	1.00	1.00
	Power given data from H_2 (smaller correlation)											
2	0.05	0.08	0.06	0.07	0.05	0.05	0.17	0.11	0.04	0.07	0.11	0.08
4	0.16	0.12	0.05	0.10	0.15	0.10	0.63	0.26	0.05	0.17	0.57	0.39
6	0.25	0.16	0.05	0.13	0.28	0.17	0.89	0.37	0.06	0.31	0.88	0.72
10	0.50	0.19	0.05	0.20	0.60	0.37	0.99	0.47	0.08	0.58	1.00	0.97
20	0.86	0.20	0.05	0.36	0.94	0.81	1.00	0.56	0.09	0.91	1.00	1.00
50	1.00	0.23	0.05	0.67	1.00	1.00	1.00	0.61	0.11	1.00	1.00	1.00
	Power given data from H_3 ($t(8)$ -distribution)											
2	0.10	0.09	0.11	0.10	0.10	0.11	0.39	0.14	0.11	0.39	0.07	0.47
4	0.22	0.09	0.11	0.19	0.08	0.28	0.88	0.14	0.12	0.80	0.06	0.94
6	0.43	0.08	0.11	0.30	0.08	0.50	0.99	0.14	0.15	0.95	0.07	1.00
10	0.79	0.09	0.11	0.52	0.08	0.84	1.00	0.14	0.14	1.00	0.06	1.00
20	0.99	0.09	0.13	0.81	0.09	1.00	1.00	0.13	0.15	1.00	0.07	1.00
50	1.00	0.08	0.13	0.97	0.08	1.00	1.00	0.13	0.15	1.00	0.06	1.00
	Power given data from H_4 (CCC-GARCH(1,1))											
2	0.26	0.20	0.43	0.28	0.44	0.31	0.39	0.27	0.48	0.43	0.48	0.46
4	0.28	0.17	0.44	0.29	0.45	0.33	0.48	0.21	0.50	0.45	0.52	0.56
6	0.32	0.15	0.45	0.30	0.47	0.37	0.56	0.20	0.48	0.49	0.52	0.65
10	0.41	0.15	0.45	0.36	0.52	0.46	0.73	0.18	0.50	0.54	0.56	0.79
20	0.54	0.14	0.46	0.40	0.58	0.57	0.90	0.18	0.50	0.62	0.60	0.92
50	0.76	0.13	0.46	0.47	0.66	0.77	1.00	0.18	0.53	0.80	0.71	1.00

Note: $Z_t^{2\ddagger}$ denotes the test proposed by [Dovern and Manner \(2020\)](#), AvR the average rank test based on [Thorarindottir et al. \(2016\)](#), ES the tests based on the energy score and LS the tests based on the log score. The subscript D refers to the entropy variant of the tests, the subscript GBT to the generalized Box transform variant. d denotes the number of variables and T the number of periods. The rejection rates are obtained using 5000 simulations, a standard t -test in the cases of ES_D and LS_D, and the raw-moments test by [Knüppel \(2015\)](#) with zero lags in all other cases. $J = J_0 = J_1 = 5000$ for all simulation-based tests.

Table 1: Size and power of tests for multivariate calibration

tation in the `bvarsv` package for R (Krüger, 2015).²⁰ Briefly, the model is given by

$$\begin{aligned} Y_t &= \nu_t + A_t Y_{t-1} + \varepsilon_t \\ \varepsilon_t &= \mathcal{N}(0, \Sigma_t), \end{aligned}$$

where Y_t is the trivariate vector covering GDP growth, inflation and the interest rate in quarter t , ε_t is a trivariate vector of error terms, and \mathcal{N} denotes the trivariate normal distribution. The model allows the parameters in ν_t , A_t and Σ_t to vary over time, following random walk type specifications. Time variation in Σ_t is especially important empirically, as it accounts for time-varying volatility in macroeconomic time series. The prior choices used for Bayesian estimation seek to balance flexibility with stability by discouraging excessive time variation in the model parameters. For comparison, we also consider a constant-parameter variant of the model, in which we choose very tight priors that ‘shut off’ time variation in η_t , A_t and Σ_t . All other prior choices are the same.

At each forecast date, we produce 100 000 draws from the model’s predictive distribution. We retain every tenth of these draws and split the resulting sequence in two halves, thus resulting in two sequences of 5000 draws each. We implement our proposed calibration tests based on these sequences, using the split-sample estimator described in Sections 3.2 and 3.3. We further obtain an analytical forecast density as required for the logarithmic score by exploiting the mixture-of-parameters form implied by the Bayesian predictive distribution (cf. Krüger et al., 2021).

Table 2 presents the results of our calibration test. The left panel refers to a sample period ending in 2019:Q4, excluding the very large Covid-related outliers at the end of the sample period.²¹ The test results provide no evidence against auto-calibration of the Primiceri model, with all p -values exceeding 5%. For the full evaluation sample including the recent pandemic period (right panel), some of the p -values of the tests based on the logarithmic score decrease, hence providing more evidence against auto-calibration. That said, the case against auto-calibration is still weak overall. The bottom panels of Table 2 refer to the constant parameter model. In both sample periods (bottom left and bottom right panels), the GBT version of the test clearly rejects auto-calibration, with all p -values being smaller than 1.2%. By contrast, the entropy version of the test does not reject auto-calibration, with all p -values exceeding 5%. We return to this discrepancy below.

The difference between realized and expected scores, which is central to the entropy variant of our test, can offer further insights on forecast calibration in applied time series settings. The top panel of Figure 2 plots this difference over time, focusing on the log score (i.e., considering $\hat{D}_{LS,t}$ defined in Equation 13) and forecast horizon $h = 0$ for brevity. The results for the energy score and for other forecast horizons are qualitatively similar. Recall that large values of $\hat{D}_{LS,t}$

²⁰A description of the package is available at https://github.com/FK83/bvarsv/blob/master/bvarsv_Nov2015_website.pdf.

²¹In particular, the annualized GDP growth rates in 2020:Q2 and 2020:Q3 equal -37.7% and +28.8%.

Stochastic volatility model (Primiceri)									
Pre-Covid sample (1965:Q4 - 2019:Q4)					Full sample (1965:Q4 - 2021:Q3)				
h	ES _D	ES _{GBT}	LS _D	LS _{GBT}	h	ES _D	ES _{GBT}	LS _D	LS _{GBT}
0	17.51	40.03	21.91	52.65	0	74.10	27.51	15.24	33.69
1	22.05	38.01	24.08	16.49	1	83.05	66.59	12.90	17.32
2	23.33	34.61	23.62	75.54	2	85.15	43.49	11.98	50.71
3	41.00	65.51	27.60	32.67	3	71.91	86.86	12.22	17.87
4	67.88	84.48	11.84	6.83	4	53.46	87.16	5.65	2.99


Constant parameter model									
Pre-Covid sample (1965:Q4 - 2019:Q4)					Full sample (1965:Q4 - 2021:Q3)				
h	ES _D	ES _{GBT}	LS _D	LS _{GBT}	h	ES _D	ES _{GBT}	LS _D	LS _{GBT}
0	5.15	0.04	88.56	0.00	0	85.48	0.03	36.83	0.00
1	13.60	0.24	97.78	0.01	1	94.41	0.16	36.90	0.01
2	20.91	0.06	97.47	0.05	2	97.77	0.05	33.72	0.02
3	30.86	0.18	94.83	0.05	3	93.13	0.23	32.50	0.02
4	48.88	1.20	75.42	0.26	4	79.67	0.67	25.70	0.18

Table 2: Test results (p -values in percent) for trivariate forecast distributions of US GDP growth, inflation, and TBILL rate. h denotes forecast horizon in quarters. Columns refer to test variants. ‘ES’ and ‘LS’ denote the energy and logarithmic score. ‘D’ and ‘GBT’ refer to entropy difference and generalized Box transform.

correspond to overconfidence (actual performance is worse than expected performance), and vice versa for small values of $\hat{D}_{LS,t}$. The plot shows no clear tendency towards either over- or underconfidence before 2020. The Covid-related outlier in 2020:Q2 corresponds to a very large value of $\hat{D}_{LS,t}$ (top right corner of the plot), indicating a massive forecast error that was essentially ruled out by the model. By contrast, the last few values for $\hat{D}_{LS,t}$ are among the smallest in the sample, indicating better-than-expected forecast performance. This result can be explained by the model’s drastic upward correction of forecast variance in response to the 2020:Q2 and 2020:Q3 observations, resulting in a forecast variance that seems implausibly high. This issue and the related discussion of how to treat Covid-19 outliers has received considerable attention recently (e.g. [Carriero et al., 2022](#)).

The bottom panel of Figure 2 shows an analogous plot for the constant parameter model. In addition to the outlier in 2020:Q2, the bottom panel features an additional outlier in 2020:Q3. This effect arises because the constant parameter model does not strongly revise its volatility prediction upward after 2020:Q2, and is hence surprised by the very high positive growth rate in 2020:Q3 reflecting a growth rebound effect. Note also that the larger value for $\hat{D}_{LS,t}$ in 2020:Q2 (about 70 in bottom panel, compared to 30 in top panel) arises because the constant parameter model features less probability mass in the far left tail of the forecast distribution, thus producing an even larger log score.

Stochastic volatility model (Primiceri)



Constant parameter model




Figure 2: Time series plot of difference between realized and expected log scores ($\hat{D}_{LS,t}$) for the macroeconomic case study at $h = 0$. Blue line shows a nonparametric time trend, based on local quadratic regression as implemented in the R function `loess`.

The bottom panel of Figure 2 also illustrates why the entropy test does not reject auto-calibration of the constant parameter model: Periods of over-confidence (such as before the great moderation of the mid-1980s) and periods of under-confidence (such as in the 1990s) offset each other, such that the hypothesis that $\mathbb{E}(\hat{D}_{LS,t}) = 0$ cannot be rejected. While this ‘blind spot’ may be viewed as a disadvantage of the entropy test, it reflects the test’s simplicity which in turn enables clear feedback for model fitting. Adjustments of the entropy test which increase power at the cost of adding complexity (e.g., an extension in the spirit of the fluctuation test of [Giacomini and Rossi, 2010](#)) are possible but are beyond the scope of this paper.

6 Finance case study: Forecasting international stock returns

As a second case study, we consider predicting the joint distribution of ten international stock indexes from Brazil (BSVP), France (FCHI), Germany (GDAXI), Hong Kong (HSI), Japan (N225), the UK (FTSE), the US (DJI), Spain (IBEX), Switzerland (SSMI), and Europe as a whole (STOXX50E). The data was kindly provided by the Oxford-Man realized library.

We consider forecasts for each business day from 2010 to 2021, based on a rolling window of

1000 business days. In addition to daily (logarithmic, close-to-close) returns, we use realized kernel measures ([Barndorff-Nielsen, Hansen, Lunde, and Shephard, 2008, 2009](#)) included in the data set. We delete days with missing return or regressor observations for either of the indexes, resulting in an evaluation sample of 2 496 observations.

Predicting high-dimensional return distributions is challenging: First, financial return data may feature stylized facts such as conditional heteroskedasticity or skewness that should be addressed for each individual index. Second, the curse of dimensionality precludes overly complex multivariate distribution models. Here we consider a pragmatic approach that seeks to address this situation:

1. For each index k , we fit a quantile regression model of the form

$$q_\tau(R_t^k|\mathcal{I}_{t-1}) = \beta_0 + \beta_1 \text{RK}_{t-1}^k + \beta_2 \times \frac{1}{5} \sum_{l=1}^5 \text{RK}_{t-l}^k + \beta_3 \times \frac{1}{22} \sum_{l=1}^{22} \text{RK}_{t-l}^k,$$

where $q_\tau(R_t^k|\mathcal{I}_{t-1})$ denotes the τ -quantile of the return R_t^k given \mathcal{I}_{t-1} , and RK_{t-l}^k denotes the realized kernel measure for index k on day $t-l$. We consider a dense grid of quantile levels $\tau = 1/1001, \dots, 1000/1001$.

2. Based on the quantile regressions, we consider a spread measure $S_t^k = \hat{q}_{.6915}(R_t^k|\mathcal{I}_{t-1}) - \hat{q}_{.3085}(R_t^k|\mathcal{I}_{t-1})$, i.e. the difference between the predicted quantiles at levels 69.15% and 30.85%. These levels are chosen such that they correspond to a spread measure of one for a standard normal distribution, thus mimicking the standard deviation.
3. For each index $k = 1, \dots, 10$ and day $t = 1, \dots, T$ in the training sample, we then compute the standardized return $Z_t^k = R_t^k/S_t^k$.
4. We then form a forecast distribution for date $T+1$ by pairing the quantile regression based univariate distributions according to the empirical (rank) copula prescribed by the standardized returns $\{Z_t^k = R_t^k/S_t^k\}_{t,k}$. This step finally yields a forecast sample of 1000 draws for the 10-variate return vector of interest.

The idea to use quantile regression at various levels for predicting a distribution has been used by [Adrian, Boyarchenko, and Giannone \(2019\)](#), among others. A distinct feature of our approach is that we consider a dense grid of 1000 quantile levels. We then sort the resulting predictions in order to obtain a monotone quantile curve (see [Chernozhukov, Fernández-Val, and Galichon, 2010](#)). Fitting a parametric curve to the quantile regression output (as in [Adrian et al., 2019](#)) is not necessary in our case since our evaluation methodology based on the energy score does not require a forecast density.

While simple to implement via the R package `quantreg` ([Koenker, 2021](#)), our approach allows for flexible univariate forecast distributions, and a flexible impact of the regressors across different parts of the distribution. Our choice of realized measures as regressors for return quantiles

follows [Žikeš and Baruník \(2015\)](#). Our use of averages over different lags (one, five and 22) follows [Corsi \(2009\)](#) who proposes this construction as a simple approximation of long memory type persistence.

Having fitted the univariate forecast distributions, the approach described in step 4 is known as the Schaake shuffle. The latter was proposed by [Clark et al. \(2004\)](#) as a pragmatic yet plausible device of pairing meteorological forecast distributions for various quantities. [Schefzik, Thorarinsdottir, and Gneiting \(2013\)](#) note its empirical copula interpretation. The key idea behind the Schaake shuffle is to draw univariate forecast distributions of the same size (1000) as the training sample, and then use the rank structure of the training sample to construct a multivariate forecast distribution. Here we use the rank structure of the standardized returns (Step 3); alternatively, the rank structure of the raw returns could be used. See [Grothe, Kächele, and Krüger \(2022\)](#) for a recent application of the Schaake shuffle to multivariate forecasting of energy prices and further discussion.

Table 3 shows our tests' results for the current case study. We focus on the energy score since a forecast density is not available in the current setup, rendering the use of the log score highly impractical. Auto-calibration of the univariate forecast distributions is not rejected at the 5% level, with a single exception (HSI, GBT test). Remarkably, auto-calibration of the ten-variate distribution for all assets is not rejected either, with both p -values exceeding 15% (last row of Table 3). As a simple check of the tests' power, we also compute p -values for a naïve multivariate forecast distribution that pairs the univariate distributions in random order (i.e., assuming independence across return indexes). Reassuringly, both test versions clearly reject auto-calibration with a very small p -value below 0.1%.

Figure 3 complements these results by plotting the entropy difference $\hat{D}_{ES,t}$ over time. The flat nonparametric time trend (blue line) indicates that there is no meaningful time variation in either under- or overconfidence, indicating that the forecast model captures the dynamics of the return distribution. Around the beginning of the Covid pandemic in March 2020, we observe a similar phenomenon as for the stochastic volatility model from Section 6: After a few extreme observations that come as a surprise to the model (large positive values for $\hat{D}_{ES,t}$), we observe several large negative values for $\hat{D}_{ES,t}$. These values can be explained by the model predicting large forecast uncertainty, in contrast to the realizing observations being moderate. However, this phenomenon is rather short-lived, and seems inevitable given the extreme (return and regressor) observations at the onset of the pandemic.

Index(es)	ES _D	ES _{GBT}
DJI	94.70	90.57
HSI	16.92	0.79
BVSP	95.65	62.61
FCHI	95.66	25.21
FTSE	45.62	57.34
IBEX	24.03	48.23
N225	33.20	33.73
SSMI	18.75	47.48
GDAXI	67.55	69.75
STOXX50E	59.48	75.91
All (ten-variate)	16.55	59.08

Table 3: Test results (p -values in percent) for the finance case study. Test variants covered: ES_D (energy score, entropy difference) and ES_{GBT} (energy score, generalized Box transform). See text for details.




Figure 3: Time series plot of difference between realized and expected energy scores ($\hat{D}_{ES,t}$) for the finance case study. Blue line shows a nonparametric time trend, based on local quadratic regression as implemented in the R function loess.

7 Conclusion

This paper proposes a framework for calibration testing of multivariate forecast distributions. This framework nests existing tests, and suggests improvements. We argue in favor of testing calibration based on proper scoring rules, and demonstrate that the resulting tests have good size and power properties in simulations. Furthermore, we show their usefulness in two challenging case studies from macroeconomics and finance. Our recommended test implementation, summarized for convenience in the following table, depends on whether power or interpretability is of higher priority to the user, and on whether the multivariate forecast density is available in closed form:

Focus on	Closed-form forecast density available?	Recommended test version	Implementation reference
Power	Yes	GBT test, log score	$\hat{U}_{LS,t}$, Eq. (11)
Power	No	GBT test, energy score	$\hat{U}_{ES,t}$, Eq. (12)
Interpretability	Yes	Entropy test, log score	$\hat{D}_{LS,t}$, Eq. (13)
Interpretability	No	Entropy test, energy score	$\hat{D}_{ES,t}$, Eq. (14)

From a general perspective, our tests compare simulated model output to realized data and reject auto-calibration in case of large discrepancies. This broad idea of model checking is widely used across the sciences. For example, posterior predictive checks (e.g. [Gabry, Simpson, Vehtari, Betancourt, and Gelman, 2019](#), Section 5) used in Bayesian statistics are based on this principle; see also [Cockayne, Graham, Oates, Sullivan, and Teymur \(2022\)](#) for a more theoretical perspective. In macroeconomics, structural economic models are often judged by how well they replicate certain features of empirical business cycle data. A common challenge across these (and many other) application areas is to decide just which ‘features of the data’ are important and should thus be replicated by the forecast model. While this question may be answered by substantive concerns (such as a specific decision problem), forecast models are often produced for multiple purposes or for communication to a broader, diverse audience. This calls for a broadly applicable strategy to summarize data. We argue that proper scoring rules are well suited for this purpose, as evidenced by their wide use in forecast comparisons.

References

- Adrian, T., N. Boyarchenko, and D. Giannone (2019). Vulnerable growth. *American Economic Review* 109, 1263–89.
- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2008). Designing realized kernels to measure the ex post variation of equity prices in the presence of noise. *Econometrica* 76, 1481–1536.
- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2009). Realized kernels in practice: Trades and quotes. *Econometrics Journal* 12, C1–C32.
- Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: A multivariate generalized ARCH model. *The Review of Economics and Statistics* 72, 498–505.
- Box, G. E. P. (1980). Sampling and Bayes’ inference in scientific modelling and robustness. *Journal of the Royal Statistical Society: Series A (General)* 143, 383–430.
- Carriero, A., T. Clark, M. Marcellino, and E. Mertens (2022). Addressing COVID-19 outliers in BVARs with stochastic volatility. *Review of Economics and Statistics*. forthcoming.
- Catania, L., S. Grassi, and F. Ravazzolo (2019). Forecasting cryptocurrencies under model and parameter instability. *International Journal of Forecasting* 35, 485–501.
- Chan, J. C. C. (2020). Large Bayesian VARs: A flexible Kronecker error covariance structure. *Journal of Business & Economic Statistics* 38, 68–79.
- Chen, Y.-T. (2011). Moment tests for density forecast evaluation in the presence of parameter estimation uncertainty. *Journal of Forecasting* 30, 409–450.
- Chernozhukov, V., I. Fernández-Val, and A. Galichon (2010). Quantile and probability curves without crossing. *Econometrica* 78, 1093–1125.
- Clark, M., S. Gangopadhyay, L. Hay, B. Rajagopalan, and R. Wilby (2004). The Schaake shuffle: A method for reconstructing space–time variability in forecasted precipitation and temperature fields. *Journal of Hydrometeorology* 5, 243–262.
- Clark, T. E. (2011). Real-time density forecasts from Bayesian vector autoregressions with stochastic volatility. *Journal of Business & Economic Statistics* 29, 327–341.
- Clements, M. P. (2014). Forecast uncertainty—ex ante and ex post: US inflation and output growth. *Journal of Business & Economic Statistics* 32, 206–216.
- Clements, M. P. and J. Smith (2002). Evaluating multivariate forecast densities: a comparison of two approaches. *International Journal of Forecasting* 18, 397 – 407.
- Cockayne, J., M. M. Graham, C. J. Oates, T. Sullivan, and O. Teymur (2022). Testing whether a learning procedure is calibrated. *Journal of Machine Learning Research* 23, 1–36.

- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* 7, 174–196.
- Craiu, R. V. and J. S. Rosenthal (2014). Bayesian computation via Markov chain Monte Carlo. *Annual Review of Statistics and Its Application* 1, 179–201.
- Dawid, A. P. (1984). Present position and potential developments: Some personal views: Statistical theory: The prequential approach. *Journal of the Royal Statistical Society: Series A (General)* 147, 278–290.
- Dawid, A. P. and P. Sebastiani (1999). Coherent dispersion criteria for optimal experimental design. *The Annals of Statistics* 27, 65 – 81.
- Del Negro, M. and G. E. Primiceri (2015). Time varying structural vector autoregressions and monetary policy: a corrigendum. *The Review of Economic Studies* 82, 1342–1345.
- Diebold, F. X., T. A. Gunther, and A. S. Tay (1998). Evaluating density forecasts with applications to financial risk management. *International Economic Review* 39, 863–883.
- Diebold, F. X., J. Hahn, and A. S. Tay (1999, November). Multivariate density forecast evaluation and calibration in financial risk management: High-frequency returns on foreign exchange. *The Review of Economics and Statistics* 81, 661–673.
- Diebold, F. X. and R. S. Mariano (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics* 13, 253–263.
- Diks, C., V. Panchenko, O. Sokolinskiy, and D. van Dijk (2014). Comparing the accuracy of multivariate density forecasts in selected regions of the copula support. *Journal of Economic Dynamics and Control* 48, 79–94.
- Dimitriadis, T., T. Gneiting, and A. I. Jordan (2021). Stable reliability diagrams for probabilistic classifiers. *Proceedings of the National Academy of Sciences* 118, e2016191118.
- Dovern, J. and H. Manner (2020). Order-invariant tests for proper calibration of multivariate density forecasts. *Journal of Applied Econometrics* 35, 440–456.
- Gabry, J., D. Simpson, A. Vehtari, M. Betancourt, and A. Gelman (2019). Visualization in Bayesian workflow. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 182, 389–402.
- Galvão, A. and J. Mitchell (2019). Measuring data uncertainty: An application using the Bank of England’s fan charts for historical GDP growth. *WBS EMF Working Paper Series no. 24*.
- Genest, C. and L.-P. Rivest (2001). On the multivariate probability integral transformation. *Statistics & Probability Letters* 53, 391–399.
- Giacomini, R. and B. Rossi (2010). Forecast comparisons in unstable environments. *Journal of Applied Econometrics* 25, 595–620.

- Giacomini, R. and H. White (2006). Tests of conditional predictive ability. *Econometrica* 74, 1545–1578.
- Gneiting, T. (2008). Editorial: Probabilistic forecasting. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 171, 319–321.
- Gneiting, T., F. Balabdaoui, and A. E. Raftery (2007). Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69, 243–268.
- Gneiting, T. and M. Katzfuss (2014). Probabilistic forecasting. *Annual Review of Statistics and Its Application* 1, 125–151.
- Gneiting, T. and A. E. Raftery (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102, 359–378.
- Gneiting, T. and R. Ranjan (2013). Combining predictive distributions. *Electronic Journal of Statistics* 7, 1747–1782.
- Gneiting, T. and J. Resin (2022). Regression diagnostics meets forecast evaluation: Conditional calibration, reliability diagrams, and coefficient of determination. *Preprint, arXiv:2108.03210*.
- Gneiting, T., L. I. Stanberry, E. P. Gruit, L. Held, and N. A. Johnson (2008). Assessing probabilistic forecasts of multivariate quantities, with an application to ensemble predictions of surface winds. *Test* 17, 211.
- González-Rivera, G. and E. Yoldas (2012). Autocontour-based evaluation of multivariate predictive densities. *International Journal of Forecasting* 28, 328 – 342.
- Grothe, O., F. Kächele, and F. Krüger (2022). From point forecasts to multivariate probabilistic forecasts: The Schaafe shuffle for day-ahead electricity price forecasting. Preprint, arxiv:2204.10154.
- Guler, K., P. T. Ng, and Z. Xiao (2017). Mincer–Zarnowitz quantile and expectile regressions for forecast evaluations under asymmetric loss functions. *Journal of Forecasting* 36, 651–679.
- Gupta, R., F. Huber, and P. Piribauer (2020). Predicting international equity returns: Evidence from time-varying parameter vector autoregressive models. *International Review of Financial Analysis* 68, 101456.
- Heinrich, C., K. H. Hellton, A. Lenkoski, and T. L. Thorarinsdottir (2021). Multivariate post-processing methods for high-dimensional seasonal weather forecasts. *Journal of the American Statistical Association* 116, 1048–1059.
- Held, L., K. Rufibach, and F. Balabdaoui (2010). A score regression approach to assess calibration of continuous probabilistic predictions. *Biometrics* 66, 1295–1305.

- Knüppel, M. (2015). Evaluating the calibration of multi-step-ahead density forecasts using raw moments. *Journal of Business & Economic Statistics* 33, 270–281.
- Ko, S. I. and S. Y. Park (2013). Multivariate density forecast evaluation: A modified approach. *International Journal of Forecasting* 29, 431 – 441.
- Koenker, R. (2021). *quantreg: Quantile Regression*. R package version 5.83.
- Koop, G. M. (2013). Forecasting with medium and large Bayesian VARS. *Journal of Applied Econometrics* 28, 177–203.
- Krüger, F. (2015). *bvarsv: Bayesian analysis of a vector autoregressive model with stochastic volatility and time-varying parameters*.
- Krüger, F., S. Lerch, T. L. Thorarinsdottir, and T. Gneiting (2021). Predictive inference based on Markov chain Monte Carlo output. *International Statistical Review* 89, 274–301.
- Krüger, F. and L. Pavlova (2022). Quantifying subjective uncertainty in survey expectations. Preprint, available at <https://ewifo.econ.kit.edu/downloads/KruegerPavlova2022.pdf>.
- Lerch, S., T. L. Thorarinsdottir, F. Ravazzolo, and T. Gneiting (2017). Forecaster’s dilemma: extreme events and forecast evaluation. *Statistical Science* 32, 106–127.
- McAlinn, K., K. A. Aastveit, J. Nakajima, and M. West (2020). Multivariate Bayesian predictive synthesis in macroeconomic forecasting. *Journal of the American Statistical Association* 115, 1092–1110.
- Mincer, J. A. and V. Zarnowitz (1969). The evaluation of economic forecasts. In J. A. Mincer (Ed.), *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*, pp. 3–46. Columbia University Press, New York.
- Mosler, K. and P. Mozharovskiy (2022). Choosing among notions of multivariate depth statistics. *Statistical Science* 37, 348–368.
- Neyman, J. (1937). ‘Smooth test’ for goodness of fit. *Scandinavian Actuarial Journal* 20, 150–199.
- Pohle, M.-O. (2020). The Murphy decomposition and the calibration-resolution principle: A new perspective on forecast evaluation. Preprint, arxiv:2005.01835.
- Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies* 72, 821–852.
- Rosenblatt, M. (1952). Remarks on a multivariate transformation. *Annals of Mathematical Statistics* 23, 470–472.

- Rossi, B. and T. Sekhposyan (2019). Alternative tests for correct specification of conditional predictive densities. *Journal of Econometrics* 208, 638–657.
- Rüschendorf, L. (2009). On the distributional transform, Sklar’s theorem, and the empirical copula process. *Journal of Statistical Planning and Inference* 139, 3921–3927.
- Schefzik, R., T. L. Thorarinsdottir, and T. Gneiting (2013). Uncertainty quantification in complex simulation models using ensemble copula coupling. *Statistical Science* 28, 616–640.
- Smith, J. Q. (1985). Diagnostic checks of non-standard time series models. *Journal of Forecasting* 4, 283–291.
- Székely, G. J. and M. L. Rizzo (2005). A new test for multivariate normality. *Journal of Multivariate Analysis* 93, 58–80.
- Theil, H. (1961). *Economic Forecasts and Policy* (2. rev. ed.). Contributions to economic analysis; 15. Amsterdam: North-Holland Publ. Co.
- Thorarinsdottir, T. L., M. Scheuerer, and C. Heinz (2016). Assessing the calibration of high-dimensional ensemble forecasts using rank histograms. *Journal of Computational and Graphical Statistics* 25, 105–122.
- Thorarinsdottir, T. L. and N. Schuhen (2018). Verification: Assessment of calibration and accuracy. In S. Vannitsem, D. S. Wilks, and J. Messner (Eds.), *Statistical postprocessing of ensemble forecasts*, pp. 155–186. Elsevier.
- Tsyplakov, A. (2011). Evaluating density forecasts: A comment. Preprint, available at https://mpra.ub.uni-muenchen.de/32728/1/MPRA_paper_32728.pdf.
- Tsyplakov, A. (2014). Theoretical guidelines for a partially informed forecast examiner. Preprint, available at <https://mpra.ub.uni-muenchen.de/67333/>.
- Vannitsem, S., D. S. H. Wilks, and J. W. Messner (Eds.) (2018). *Statistical Postprocessing of Ensemble Forecasts*. Amsterdam, Netherlands: Elsevier.
- Wei, W., F. Balabdaoui, and L. Held (2017). Calibration tests for multivariate Gaussian forecasts. *Journal of Multivariate Analysis* 154, 216–233.
- West, K. D. (1996). Asymptotic inference about predictive ability. *Econometrica* 64, 1067–1084.
- Wilks, D. S. (2017). On assessing calibration of multivariate ensemble forecasts. *Quarterly Journal of the Royal Meteorological Society* 143, 164–172.
- Ziegel, J. F. (2015). Copula calibration. *Oberwolfach Reports* 12, 1091–1094.
- Ziegel, J. F. and T. Gneiting (2014). Copula calibration. *Electronic Journal of Statistics* 8, 2619–2638.

- Ziel, F. and K. Berk (2019). Multivariate forecasting evaluation: On sensitive and strictly proper scoring rules. Preprint, arxiv:1910.07325.pdf.
- Žikeš, F. and J. Baruník (2015). Semi-parametric conditional quantile models for financial returns and realized volatility. *Journal of Financial Econometrics* 14, 185–226.
- Zuo, Y. and R. Serfling (2000). General notions of statistical depth function. *Annals of Statistics* 28, 461–482.

A Additional simulations

A.1 Size

Here we explore the size of our proposed methods in two setups that are potentially more challenging than the one in Section 4: First, a setup involving time series dependence in forecasts and realizations. Second, a setup in which the true model follows a multivariate t (rather than normal) distribution.

Time series dependence

We simulate data from a d -dimensional vector autoregressive (VAR) model of order one:

$$Y_t = A Y_{t-1} + \varepsilon_t,$$

where Y_t and ε_t are $d \times 1$ vectors of random variables with $d \in \{2, 4, 10\}$. All diagonal elements of the matrix A are equal to 0.5, and all off-diagonal elements are equal to zero. The error term vector ε_t follows a multivariate normal distribution, with mean vector zero and covariance matrix $\gamma_t \Sigma$, with Σ a $d \times d$ matrix and γ_t being a scalar factor. Σ has an equicorrelation structure with all diagonal elements equal to one and all off-diagonal elements equal to 0.5. In the homoskedastic case, $\gamma_t = 1$ for all t ; in the heteroskedastic case, $\gamma_t \in \{1, 1.25\}$ follows a first-order Markov Chain with two states. In each state, the probabilities of remaining and leaving are equal to 0.7 and 0.3 respectively. We consider forecast horizons of one and four periods, with the latter choice implying a three-period overlap in subsequent forecast errors. In order to account for possible autocorrelation under the null hypothesis, we use HAC approaches for implementing the t -test (used for the entropy tests) and the Knüppel (2015) uniformity test (used for the GBT tests).²²

Table 4 presents simulation results for a sample of size $T = 200$. The rejection rates of all four test variants are between 0.04 and 0.08, and thus close to the nominal rate of 0.05. These results indicate that time series dependence and the associated use of HAC approaches pose no severe challenges in the present setup.

t distribution

We additionally consider a static setup where Y_t follows a d -variate t distribution with eight degrees of freedom. We consider $d \in \{2, 4, 10\}$ and a sample size of $T = 200$. Table 5 shows that the rejection rates are between 0.05 and 0.06 in all cases, again indicating good size properties.

²²Our implementation of the uniformity test closely follows the test statistic called α_{1234}^0 in Knüppel (2015). For the robust t -test, we use a HAC covariance estimator analogous to the one by Knüppel (2015).

d	Horizon	Heterosk.?	ES _D	ES _{GBT}	LS _D	LS _{GBT}
2	1	No	0.05	0.04	0.06	0.05
4	1	No	0.06	0.05	0.05	0.05
10	1	No	0.06	0.05	0.05	0.05
2	4	No	0.07	0.06	0.08	0.05
4	4	No	0.07	0.05	0.07	0.05
10	4	No	0.07	0.05	0.07	0.05
2	1	Yes	0.06	0.05	0.05	0.05
4	1	Yes	0.06	0.05	0.06	0.05
10	1	Yes	0.05	0.04	0.05	0.04
2	4	Yes	0.07	0.05	0.07	0.05
4	4	Yes	0.06	0.05	0.07	0.05
10	4	Yes	0.08	0.05	0.07	0.05

Table 4: Size under time series dependence with $T = 200$. Nominal size equals 0.05. Rejection rates are based on 5000 simulations.

d	ES _D	ES _{GBT}	LS _D	LS _{GBT}
2	0.06	0.05	0.06	0.05
4	0.06	0.06	0.06	0.05
10	0.05	0.05	0.05	0.05

Table 5: Size in a t -distributed setup (10 degrees of freedom) with $T = 200$. Nominal size equals 0.05. Rejection rates are based on 5000 simulations.

T	h	Heterosk.?		Σ		A		Rejection Rates			
		True	Fcst	True	Fcst	True	Fcst	ES _D	ES _{GBT}	LS _D	LS _{GBT}
50	1	Yes	Yes	1	1.0	0.5	0.8	0.13	0.05	0.17	0.06
200	1	Yes	Yes	1	1.0	0.5	0.8	0.39	0.23	0.55	0.33
50	4	Yes	Yes	1	1.0	0.5	0.8	0.14	0.05	0.17	0.05
200	4	Yes	Yes	1	1.0	0.5	0.8	0.33	0.20	0.45	0.27
50	1	Yes	Yes	1	1.2	0.5	0.5	0.46	0.14	0.82	0.25
200	1	Yes	Yes	1	1.2	0.5	0.5	0.91	0.82	1.00	0.99
50	4	Yes	Yes	1	1.2	0.5	0.5	0.41	0.09	0.68	0.14
200	4	Yes	Yes	1	1.2	0.5	0.5	0.76	0.61	0.99	0.88
50	1	Yes	No	1	1.0	0.5	0.5	0.12	0.05	0.17	0.06
200	1	Yes	No	1	1.0	0.5	0.5	0.42	0.25	0.57	0.35
50	4	Yes	No	1	1.0	0.5	0.5	0.11	0.04	0.14	0.04
200	4	Yes	No	1	1.0	0.5	0.5	0.29	0.16	0.42	0.25

Table 6: Power under time series dependence. Here auto-calibration is violated, and rejection rates represent power. Rows 1-4 correspond to misspecified persistence (diagonal elements of A equal to 0.8 instead of 0.5). Rows 5-8 correspond to misspecified residual variance (diagonal elements of Σ equal to 1.2 instead of 1.0). Rows 9-12 correspond to neglected heteroskedasticity. Rejection rates are based on 5000 simulations.

A.2 Additional simulations: Power

Table 6 presents further simulation results on the power of the tests. While the true data is simulated from the same VAR model as in Section A.1, we consider various types of deviations of the forecast model from the true model. The left half of Table 6 lists the specific setups considered. As expected, all tests have higher power for the larger sample size $T = 200$, as compared to $T = 50$. Furthermore, we generally observe a slight decrease in power when moving to a higher forecast horizon ($h = 4$ instead of $h = 1$), which seems plausible given the overlap and increased persistence in the relevant calibration statistics. As in the baseline simulations of Section 4, the tests based on the log score have somewhat higher power than the ones using the energy score. The entropy variant of the test tends to have higher power than the GBT variant in the dynamic settings considered here, which is likely due to fact that all settings imply wrong variance forecasts, corresponding to setting H_1 in Table 1.