

Interpretation of point forecasts with unknown directive

Patrick Schmidt[†]

*Heidelberg Institute for Theoretical Studies, Heidelberg, Germany and
Goethe University Frankfurt, Frankfurt, Germany*

Matthias Katzfuss

Texas A&M University, College Station, USA

Abstract.

Point forecasts can be interpreted as functionals (i.e., point summaries) of predictive distributions. We consider the situation of unknown directives and show how to identify the functional based on point forecasts and associated realizations. Focusing on state-dependent quantiles and expectiles, we construct a generalized method of moments estimator for the functional, along with tests of optimality and more specific hypotheses. Using simulations, we demonstrate that our optimality test is better calibrated and more powerful than existing solutions. In a data example, we show that the gross domestic product (GDP) Greenbook forecasts of the U.S. Federal Reserve can be interpreted as state-dependent quantiles.

Keywords: functional, expectile, forecast optimality, generalized method of moments, quantile, unknown loss function

1. Introduction

Forecasts are frequently the basis of crucial decisions. Yet, they are fraught with uncertainty due to imperfections in the observation, understanding, and modeling of the underlying mechanisms. To account for this uncertainty, it is increasingly recognized that forecasts should be probabilistic in nature. If forecasts are issued in the form of predictive distributions, it is straightforward to compute the action that maximizes the expected utility, test for optimality, or compare them to other forecasts (see Gneiting and Katzfuss, 2014, for a recent review of probabilistic forecasting).

However, point forecasts are still ubiquitous. Their interpretation requires assumptions on the decision process or directive that the forecasters used to generate the point forecasts (Elliott and Timmermann, 2008; Engelberg et al., 2009; Manski, 2016). A directive can be expressed through a functional (i.e., a real-valued summary) of the predictive distribution. It is a widely used assumption that the reported functional is the mean. However, there is often little justification for this choice. Knowing which functional was used by the forecaster is important, as it allows for proper interpretation, evaluation, testing, and comparison of point forecasts (Gneiting, 2011).

We consider point forecasts with an unknown directive, for which the forecaster only implicitly reported a certain functional of the predictive distribution. This situation can

[†]*Address for correspondance:* Patrick Schmidt, HITS gGmbH, Schloss-Wolfsbrunnenweg 35, 69118 Heidelberg, Germany. E-mail: Patrick.Schmidt@h-its.org.

arise with expert forecasts or response items in surveys. Another important example for forecasts with unknown directives is output from complex computer models, which are often tuned by multiple individuals to achieve forecasts that the individuals perceive as optimal in a way that might neither be transparent nor explicitly defined. Such forecasts would be most informative if the user knew the directive under which the forecast was issued. Our goal here is to estimate the functional from a time series of point forecasts and associated realizations, and to construct tests regarding the properties of the functional. Once the functional has been estimated, the point forecasts can be coherently interpreted, improved, and compared to other point or probability forecasts.

Past work on estimating a directive based on point forecasts and realizations has focused on estimation of the loss function. Elliott et al. (2005) provide a generalized method of moments (GMM) estimator of the loss function for constant preferences and linear forecasting models. Patton and Timmermann (2007) apply this method to the U.S. Federal Reserve’s gross domestic product (GDP) forecasts with a new class of loss functions, which consists of quadratic splines that are flexible with respect to a state variable. Recently, piecewise linear and piecewise quadratic loss functions have been used in various economic applications (Christodoulakis and Mamatzakis, 2008; Capistrán, 2008; Elliott et al., 2008; Krol, 2013; Pierdzioch et al., 2013; Wang and Lee, 2014; Fritsche et al., 2015). Komunjer and Owyang (2012) derive related estimators for multivariate forecasts and loss functions. Lieli and Stinchcombe (2013) discuss the recoverability of the loss function if conditional distributions are observable. In a neuroscience application, Körding and Wolpert (2004) estimate the loss function implicit in human sensorimotor control by varying targets in an experimental task. Sims (2015) uses a similar approach to infer the implicit loss function of the visual working memory. Guler et al. (2017) propose Mincer-Zarnowitz quantile and expectile regressions to account for asymmetric loss functions.

Here, we argue that the loss function is, in fact, not identifiable based solely on point forecasts and realizations. Hence, we propose to formalize point forecasts via functionals rather than loss functions. This allows for a more general definition of forecast optimality, and we show the existence of identifying moment conditions under weaker conditions. For estimation, we focus on state-dependent quantiles and expectiles, for which the level of asymmetry can depend on the current state. We propose a GMM estimator and show consistency and asymptotic normality under mild assumptions. We also discuss testing of forecast optimality and other forecast properties.

Our approach generalizes the findings of Elliott et al. (2005) to state-dependent forecasts that are not required to be linear functions of the instrumental variables. In comparison to Patton and Timmermann (2007), who consider state-dependent loss functions, our methods are more interpretable and fully theoretically justified. Further, our approach can be used for performance comparisons between point and probability forecasts and for creating density forecasts from point forecasts.

In a data example we illustrate that our approach yields accessible and scientifically relevant insights. Specifically, we show that the GDP Greenbook forecasts of the U.S. Federal Reserve can be interpreted as state-dependent quantiles. In a Monte Carlo study, our approach exhibits better calibrated and more powerful optimality tests than existing solutions, and we demonstrate that our more general definition of optimality can be used

to distinguish different state-dependent forecasting behavior.

This manuscript is organized as follows. In Section 2, we introduce optimal forecasts and discuss non-parametric identification. In Section 3, we introduce a parametric GMM estimator in the time series setting, study its large sample behavior, and discuss tests of optimality and more specific hypotheses. In Section 4, we apply the method to the GDP Greenbook forecasts. In Section 5, we compare our methodology to the approach of Patton and Timmermann (2007) using simulated data. Section 6 serves as a discussion. Technical results and proofs are provided in the Appendix. Sections S1 through S4 contain additional details in an online Supplementary Material document.

2. Identification

In this section, we discuss non-parametric identification of the functional. Then, we describe the relationship between functionals and loss functions, and argue that loss functions cannot be identified.

Consider a real-valued random variable Y and a corresponding point forecast X , which is based on the information available to the forecaster, as encoded by some σ -algebra \mathcal{F} . Commonly, a point forecast is interpreted as the mean of the conditional distribution $\mathcal{L}(Y|\mathcal{F})$, i.e.,

$$X = \mathbb{E}[Y|\mathcal{F}].$$

Here and throughout the paper, equality of random variables is understood to hold almost surely. We proceed to a more general framework. Let $\alpha : \mathcal{P} \mapsto \mathbb{R}$ be a functional (Horowitz and Manski, 2006; Huber and Ronchetti, 2009, p. 9), i.e., a single-valued mapping from some class of probability distributions to the real line. We use the short notation $\alpha(Y|\mathcal{F})$ for $\alpha(\mathcal{L}(Y|\mathcal{F}))$.

DEFINITION 1 (OPTIMAL α -FORECAST). *A random variable X is an optimal α -forecast of Y with respect to the information set \mathcal{F} if*

$$X = \alpha(Y|\mathcal{F}).$$

Throughout, we call a functional α symmetric if, for every symmetric distribution P with symmetry point c , it holds that $c = \alpha(P)$. Prominent alternatives to the mean functional are symmetric functionals, like the median, or asymmetric generalizations such as quantiles and expectiles.

Now, crucially, we consider the situation in which the functional used by the forecaster *and* the conditional distributions $\mathcal{L}(Y|\mathcal{F})$ are unknown. In line with seminal extant work on professional economic forecasters (Elliott et al., 2005; Patton and Timmermann, 2007), we merely assume that the unknown conditional distribution constitutes a predictive distribution consistent with some information set \mathcal{F} .

We consider single-valued scalar functionals throughout, although the results extend to the set-valued case under additional technical considerations. If R is a random variable (or vector), the relation $R \in \mathcal{F}$ indicates that R is \mathcal{F} -measurable. The partial derivative of a function $g(x, y)$ with respect to x is denoted as $g_{(x)}(x, y)$.

2.1. Identifying moment conditions for functionals

It is well-known that an optimal mean-forecast relative to the information set \mathcal{F} implies an identifying moment condition (Diebold and Lopez, 1996):

$$X = \mathbb{E}[Y|\mathcal{F}] \iff \mathbb{E}[(Y - X)W] = 0 \text{ for all } W \in \mathcal{F}, \quad (1)$$

where the components of the random vector W (henceforth called *instruments*) honor the information \mathcal{F} available to the forecaster when the prediction is issued. The relation (1) allows testing whether a point forecast is an optimal mean-forecast. In the hypothetical limit of an infinite supply of data and instruments, a non-rejection of the test is a sufficient condition for mean-forecast optimality.

This property of optimal mean-forecasts can be generalized to optimal α -forecasts: Let X be a forecast based on some information set \mathcal{F} . For every sufficiently regular functional $\alpha : \mathcal{P} \mapsto \mathbb{R}$, there exists a function V identifying the optimal α -forecast, i.e.,

$$X = \alpha(Y|\mathcal{F}) \iff \mathbb{E}[V(X, Y)W] = 0 \text{ for all } W \in \mathcal{F}. \quad (2)$$

For a formal statement and regularity conditions see Lemma 1 in Appendix A. The proof applies recent results on identification functions (Steinwart et al., 2014) in the prediction space setting (Gneiting et al., 2007; Gneiting and Ranjan, 2013; Strähl and Ziegel, 2017). We can find the identification function V of a functional α via loss functions as described in Section 2.2 or by elementary considerations as exemplified in Section 3.1. The regularity conditions on the functional in Lemma 1 exclude some functionals (e.g., the mode) for their lack of continuity, and others (e.g., the variance) because they do not induce convex level sets on the class of absolutely continuous distributions.

The moment conditions in (2) identify the functional only on the class of arising conditional distributions $\mathcal{L}(Y|\mathcal{F})$. For example, the mean and the median are distinct functionals in general, but if the considered distribution is symmetric, the two functionals are identical and cannot be identified. In Section 3, we focus on optimal forecasts in the form of either state-dependent quantiles or state-dependent expectiles, which allows for unique identification without involving unduly strong assumptions on the data-generating process.

2.2. Relationship between functionals and loss functions

Previous work (e.g., Elliott et al., 2005) defined optimal point forecasts via loss functions $L(x, y)$ as

$$X = \arg \min_{x \in \mathbb{R}} \mathbb{E}[L(x, Y)|\mathcal{F}]. \quad (3)$$

Under regularity conditions, Equation (3) specifies a well-defined optimal α_L -forecast, where

$$\alpha_L : \mathcal{P} \mapsto \mathbb{R}, \quad P \mapsto \arg \min_{x \in \mathbb{R}} \mathbb{E}_{Y \sim P}[L(x, Y)]$$

for a suitable class \mathcal{P} of probability distributions. For example, the mean-functional can be defined as minimizing expected quadratic loss, $L(x, y) = (x - y)^2$, for probability distributions with finite second moments. Some functionals, such as the expected shortfall and the mode, cannot be defined via loss functions for broad classes of probability distributions (Gneiting, 2011; Heinrich, 2014). Consequently, the optimal α -forecast

from Definition 1 constitutes a more general optimality. The conditions for the existence of the identifying moment conditions (2) are satisfied for any functional defined via a continuous loss function; see Appendix A for details. In fact, every functional α with identifying moment conditions (2) can be defined via a loss function under regularity conditions on $\mathcal{L}(Y|\mathcal{F})$. In this case, the identification function V derives from the partial derivative $L_{(x)}(x, y)$ of any loss function L that defines α (Steinwart et al., 2014, Thm. 8; Fissler and Ziegel, 2016, Thm. 3.2). This fact can be used to find identification functions for functionals defined via loss functions.

2.3. Nonidentifiability of loss functions

For a specific functional, there might be many loss functions defining it (Gneiting, 2011; Ehm et al., 2016). It is therefore impossible to identify the shape of the loss, as all these functions lead to the same functional-forecast and identical moment conditions. For example, given any convex and differentiable function Φ , the Bregman loss function, $L(x, y) = \Phi(y) - \Phi(x) - \Phi'(x)(y - x)$, induces an optimal mean-forecast (Savage, 1971). Hence, loss functions are not identified, but functionals are identified to the extent that they differ on the predictive distributions.

3. Parametric estimation and testing of state-dependent quantiles and expectiles

We turn to parametric estimation of possibly varying functionals in the time series setting. Consider a stochastic process $\{(X_t, Y_t, Z_t) : t = 1, 2, \dots\}$ of forecasts, observations, and state variables, for which we have a sample path $\{(x_t, y_t, z_t) : t = 1, \dots, T\}$. Our goal is to infer the functional that the point forecasts represent.

We assume that at each point in time an optimal α -forecast is issued, i.e.,

$$X_t = \alpha(Y_t|\mathcal{F}_t) \text{ for } t = 1, 2, \dots$$

In the situation of an h -step ahead forecast, the available information is typically generated by lagged variables of the outcome Y and the vector-valued state variable Z , in which case $\mathcal{F}_t = \sigma(Y_1, \dots, Y_{t-h}, Z_1, \dots, Z_{t-h})$. For ease of notation, statements about all time points are often denoted without subscripts. For example, we write $X = \alpha(Y|\mathcal{F})$ instead of $X_t = \alpha(Y_t|\mathcal{F}_t)$ for $t = 1, 2, \dots$

Extending Definition 1, we allow the functional α to depend on the current situation, represented by the \mathcal{F} -measurable state variable Z . We call this a *state-dependent functional*. Asymmetric and state-dependent point forecasts can arise for a variety of reasons, including varying preferences of the forecaster, asymmetric information, and non-linear transformation of data (see Section S1 for details).

In light of the results and discussion in Section 2, we assume that the true functional is a state-dependent quantile, or that it is a state-dependent expectile of $\mathcal{L}(Y|\mathcal{F})$. By restricting the class of feasible functionals to only quantiles (or expectiles), the functionals induce distinct forecasts under minimal assumptions, which guarantees unique identification even under state-dependence.

3.1. State-dependent quantiles and expectiles

The τ -quantile functional $q_\tau(P)$ of a distribution P with continuous and strictly increasing cumulative distribution function is the unique solution x to the equation $P((-\infty, x]) = \tau$. We can express this directly in terms of the identification function of the τ -quantile, namely $V_{q_\tau}(x, y) = \mathbb{1}(y \leq x) - \tau$:

$$X = q_\tau(Y|\mathcal{F}) \iff \mathbb{E}[(\mathbb{1}(Y \leq X) - \tau)W] = 0 \text{ for all } W \in \mathcal{F}.$$

While quantiles are asymmetric generalizations of the median, expectiles are analogously defined as asymmetric generalizations of the mean. Specifically, the τ -expectile $e_\tau(P)$ of a distribution P with finite mean was introduced in Newey and Powell (1987) as the unique solution x to the equation

$$\frac{\tau}{1 - \tau} = \frac{\int_{-\infty}^x (x - y)dP(y)}{\int_x^{\infty} (y - x)dP(y)}.$$

This is equivalent to

$$X = e_\tau(Y|\mathcal{F}) \iff \mathbb{E}[|\mathbb{1}(Y \leq X) - \tau|(X - Y)W] = 0 \text{ for all } W \in \mathcal{F},$$

which reveals the corresponding identification function, namely $V_{e_\tau}(x, y) = |\mathbb{1}(x \geq y) - \tau|(x - y)$.

Hence, quantiles and expectiles can be identified under weak assumptions on the conditional distributions. We allow for additional flexibility and let the level τ of the quantile or expectile depend on the state variable z via a parametric function $m(z, \theta)$.

DEFINITION 2 (SPECIFICATION MODEL). *Let Θ be a subset of \mathbb{R}^p and suppose that the state variable z takes values in \mathbb{R}^k . A specification model is a function $m(z, \theta)$ that maps $\mathbb{R}^k \times \Theta$ into the unit interval $(0, 1)$.*

We say that a specification model $m(z, \theta)$ is continuous (continuously differentiable) if it is continuous (continuously differentiable) in $\theta \in \Theta$ for every z .

Let us consider examples for such specification models in Table 1, where we assume that z is real-valued. The special case of a *constant* model assumes that the forecaster always states the θ -quantile or expectile and was implemented in previous work (Elliott et al., 2005; Christodoulakis and Mamatzakis, 2008; Krol, 2013; Pierdzioch et al., 2013; Fritsche et al., 2015). The *break* model generalizes the constant model to allow for a structural break in the risk assessment at the threshold value t . The *linear* model specifies the dependence of the quantile or expectile level on the state z , where the logistic function $\Psi(x) = (1 + \exp(-x))^{-1}$ ensures that it lies in the unit interval. Lastly, we suggest a method to detect periodic asymmetry in the forecasts. The *periodic* model provides information about the base level θ_1 and the magnitude θ_2 and period θ_3 of the periodic asymmetry.

3.2. GMM estimator

Let us now assume that the forecast X is an optimal quantile-forecast with state-dependent level described by the specification model $m(z, \theta_0)$,

$$X_t = q_{m(Z_t, \theta_0)}(Y_t|\mathcal{F}_t) \quad \text{for } t = 1, 2, \dots \quad (4)$$

Table 1. Specification models.

Name	Model	Θ
Constant	$m(z, \theta) = \theta$	$(0, 1)$
Break	$m(z, \theta) = \mathbb{1}(z \leq t)\theta_1 + \mathbb{1}(z > t)\theta_2$	$(0, 1)^2$
Linear	$m(z, \theta) = \Psi(\theta_1 + z\theta_2)$	\mathbb{R}^2
Periodic	$m(z, \theta) = \Psi(\theta_1 + \theta_2 \sin(2\pi z/\theta_3))$	$\mathbb{R}^2 \times \mathbb{R}^+$

Crucially, we assume that Z_t is \mathcal{F}_t -measurable and it follows that

$$\mathbb{E}[(\mathbb{1}(Y_t \leq X_t) - m(Z_t, \theta_0)) W_t] = 0 \text{ for all } W_t \in \mathcal{F}_t \text{ for } t = 1, 2, \dots$$

We refer to $g(\theta) = (\mathbb{1}(y \leq x) - m(z, \theta)) w$ as the *moment function*. Given a sample path of \mathcal{F}_t -measurable instrumental variables $w_t = (w_{t,1}, \dots, w_{t,q})'$, the empirical mean of the moment function is given by

$$g_T(\theta) = \frac{1}{T} \sum_{t=1}^T (\mathbb{1}(y_t \leq x_t) - m(z_t, \theta)) w_t. \quad (5)$$

Then, the standard GMM estimator is obtained by minimizing the quadratic norm of the empirical moment:

$$\hat{\theta}_T = \arg \min_{\theta \in \Theta} g_T(\theta)' \hat{S}_T^{-1} g_T(\theta), \quad (6)$$

where the weighting matrix derives from a heteroskedasticity and auto-correlation consistent (HAC) estimator \hat{S}_T of the covariance matrix of the moment function g , as proposed by Newey and West (1987). Throughout this study, we use the standard two-step GMM procedure proposed in Hansen (1982) to find \hat{S}_T^{-1} in (6).

If the specification model m is continuous, there exists an \mathcal{F}_t -measurable instrumental variable W_t such that the GMM estimator is consistent for the true parameter value $\theta_0 \in \Theta$. For a formal statement see Theorem 1 in Appendix B. The crucial point is to establish unique identification of the parameter without restrictive assumptions on the unobservable conditional distributions $\mathcal{L}(Y|\mathcal{F})$. Analogously, one can assume an expectile forecast $X = e_{m(Z, \theta_0)}(Y|\mathcal{F})$ and use the moment function $g(\theta) = (\mathbb{1}(y \leq x) - m(z, \theta))(x - y)w$, which also leads to a consistent estimator.

We thus establish the existence of a suitable instrument to achieve consistency without assuming a linear forecasting model nor knowledge of the forecasting model's parameters. Note that under the assumption of a linear forecasting model, Elliott et al. (2005) obtain the stronger result that any subset of the used information identifies their one-dimensional parameter. We discuss the choice of the instrumental variables w in Section 3.3.

Once the identification of the system is established, standard GMM theory (Hansen, 1982) provides a range of useful asymptotic results. If the specification model $m(z, \theta)$ is continuously differentiable, the GMM estimator is asymptotically normal,

$$\sqrt{T}(\hat{\theta}_T - \theta_0) \rightarrow \mathcal{N}_p(0, (G' S^{-1} G)^{-1}), \text{ as } T \rightarrow \infty, \quad (7)$$

where p is the dimension of the parameter vector, S is the covariance matrix of the moment function, and G is the expectation of its partial derivative with respect to

θ . We have $G = \mathbb{E}[m_{(\theta)}(Z, \theta_0)W]$ and $S = \mathbb{E}[m(Z, \theta_0)^2 W'W]$ for quantiles, and $G = \mathbb{E}[m_{(\theta)}(Z, \theta_0)(X - Y)W]$ and $S = \mathbb{E}[m(Z, \theta_0)^2 (X - Y)^2 W'W]$ for expectiles. See Supplementary Section S2 for details.

3.3. Testing optimality with unknown directive

The so-called test of overidentifying restrictions (Hansen, 1982) can be used to test forecast optimality. Specifically, if the dimension q of the instrument vector W is greater than the dimension p of the parameter vector θ , it holds that

$$J_T(\hat{\theta}_T) \rightarrow \chi_{q-p}^2, \text{ as } T \rightarrow \infty,$$

where $J_T(\theta) = Tg_T(\theta)'\hat{S}_T^{-1}g_T(\theta)$ is called the J -statistic.

We now discuss an important aspect of our optimality definition: A point forecast can only be defined as optimal with respect to a specific functional *and a specific information set*. The choice of instruments W in (5) determines the information set for which we test. If a forecast is optimal with respect to \mathcal{F} , it also satisfies the moment conditions for any information set $\mathcal{G} \subseteq \mathcal{F}$, because

$$X = \alpha(Y|\mathcal{F}) \iff \mathbb{E}[V(X, Y)|\mathcal{F}] = 0 \Rightarrow \mathbb{E}[V(X, Y)|\mathcal{G}] = 0,$$

where V is the identification function of α . If a test with instruments W rejects optimality, the point forecast is not optimal with respect to any information set \mathcal{F} that contains the information set $\sigma(W)$ generated by W :

$$\mathbb{E}[V(X, Y)|\sigma(W)] \neq 0 \Rightarrow X \neq \alpha(Y|\mathcal{F}).$$

Thus, the null hypothesis in the test of overidentifying restrictions is

$$H_0 : \text{There exists } \theta_0 \in \Theta \text{ such that } X = q_{m(Z, \theta_0)}(Y|\mathcal{F}) \text{ with } \sigma(W) \subseteq \mathcal{F}.$$

We use this in Section 5.1 to explore the size and power of optimality tests.

An optimal yet uninformed point forecast can only be rejected if appropriate instruments are available. Furthermore, a misspecified or non-optimal forecast can still form an optimal forecast with respect to a smaller information set or a more flexible class of functionals. For this reason, the choice of instruments and specification models is a crucial part of inference based on our estimators and optimality tests. To obtain power against forecasts that are optimal with respect to different specification models, we propose to include the forecast as an instrument; see Section S4.

3.4. Specification tests for forecasting behavior

The estimation of specification models $m(z, \theta)$ provides a suitable framework for testing specific hypotheses about forecasting behavior. In general, any restriction $R(\theta_0) = 0$ for the model $m(z, \theta)$, where $R : \Theta \mapsto \mathbb{R}^l$ is differentiable, can be tested using a Wald statistic (e.g., Greene, 2012) of the form

$$W_T(\theta) = T \cdot R(\theta)'(R_{(\theta)}(\theta)\hat{S}_T(\theta)R_{(\theta)}(\theta)')^{-1}R(\theta),$$

where $R_{(\theta)}(\theta) = \partial R(\theta)/\partial\theta$ is the gradient of R evaluated at θ , and $R_{(\theta)}(\theta_0)\hat{S}_T(\theta_0)R_{(\theta)}(\theta_0)'$ is non-degenerate. If $R(\theta_0) = 0$, then

$$W_T(\hat{\theta}_T) \rightarrow \chi^2_l, \text{ as } T \rightarrow \infty. \quad (8)$$

For illustration, we discuss restrictions in the models introduced in Table 1. In the linear model, we can test if the forecasting behavior is constant with respect to the state z , which corresponds to the restriction $\theta_2 = 0$. The break model facilitates tests for the hypothesis that there is no structural break, in which case $\theta_1 - \theta_2 = 0$. This specification test is closely related to the state-dependence test in Caunedo et al. (2013). Finally, we can use the periodic model to test for the presence of periodic asymmetry, i.e., $\theta_2 = 0$. In all of these examples, joint tests of optimality and of the restriction can be carried out based on the asymptotic distribution in (8).

4. GDP growth forecasts as state-dependent quantiles

We present an application to the GDP Greenbook forecasts of the Federal Reserve. As realized values we take the quarterly real GDP growth rate over the period 1969 to 2011 ($T = 172$ observations) as reported in the first release[†]. Standard tests based on the mean functional reject optimality. Patton and Timmermann (2007) model the loss function as a quadratic spline with three nodes whose shape is allowed to change with the growth rate.

Here, we interpret the forecasts as state-dependent quantiles of the Federal Reserve’s (implicit) predictive distributions. To investigate whether the reported quantile changes with the predicted GDP growth rate, we apply the linear specification model $m(x, \theta) = \Psi(\theta_1 + x\theta_2)$ as described in Section 3.1. As instrumental variables w for the GMM estimator, we use a constant, the forecast and the one-quarter-lagged value of the outcome. In a test of overidentifying restrictions (Section 3.3), we obtain a J -statistic of 1.63 with a p -value of 0.20. Consequently, there is no reason to reject optimality if we allow for state-dependent quantile forecasts.

Compared to the spline loss function of Patton and Timmermann (2007), we use two instead of six parameters, and our more powerful test (Section 5.1) does not reject the hypothesis of an optimal forecast. The need for additional instruments due to the large number of parameters in the spline approach is a concern for small sample sizes.

We obtain the estimate $\hat{\theta}_T = (-0.20, 0.18)'$. As illustrated in Figure 1, the forecasts can be interpreted as $m(x, \hat{\theta}_T)$ -quantiles that depend on the predicted growth rate x . The Federal Reserve reports lower quantile levels during times of low growth, so forecasts are more conservative during recessions. The covariance estimate implied by (7) is given by

$$\frac{1}{T}(G_T'\hat{S}_T^{-1}G_T)^{-1} = \begin{pmatrix} \hat{\sigma}_{1,1} & \hat{\sigma}_{1,2} \\ \hat{\sigma}_{1,2} & \hat{\sigma}_{2,2} \end{pmatrix} = \begin{pmatrix} 0.082 & -0.015 \\ -0.015 & 0.005 \end{pmatrix},$$

where G_T is the sample moment of G evaluated at $\theta = \hat{\theta}_T$. As $\Psi(\cdot)$ is strictly monotone, we can compute pointwise confidence intervals for $\hat{\theta}_1 + x\hat{\theta}_2$ and transform into confidence intervals for $m(x, \hat{\theta}) = \Psi(\hat{\theta}_1 + x\hat{\theta}_2)$, as illustrated in Figure 1. The Wald test as

[†]The results in this section are robust to using the second revision or the 2017Q1 vintage.




Figure 1. The estimated specification model $m(x, \hat{\theta}_T)$ plotted against the predicted growth rate x , with pointwise confidence intervals at levels 0.6 and 0.9.

introduced in Section 3.4 for the hypothesis that $\theta_2 = 0$ (p -value = 0.01) suggests that the underlying preferences are not only asymmetric but also flexible with respect to a state variable.

When applying the lagged outcome y_{t-1} instead of the predicted value x_t as state variable the test of overidentifying restrictions rejects optimality. Consequently, the data is not consistent with a forecaster simply overweighting the predictive content of the current growth rate.

5. Simulation study

In this section, we demonstrate that our approach, while flexible, produces well calibrated and powerful optimality tests. For ease of comparison to the related approach of Patton and Timmermann (2007), we adopt their simulation setting. Each sample path y_1, \dots, y_T is simulated from an AR(1)-GARCH(1,1)-model of the form

$$\begin{aligned} Y_t &= \frac{1}{2}Y_{t-1} + \sigma_t\epsilon_t \quad \text{for } t = 1, 2, \dots, T, \\ \sigma_t^2 &= \frac{1}{10} + \frac{4}{5}\sigma_{t-1}^2 + \frac{1}{10}\sigma_{t-1}^2\epsilon_{t-1}^2, \\ \epsilon_t &\stackrel{iid}{\sim} \mathcal{N}(0, 1). \end{aligned} \tag{9}$$

We generate 2,000 paths for seven sample sizes $T \in \{50, 100, 250, 500, 1000, 2000, 4000\}$.

5.1. Constant expectile-forecasts with varying information sets

We generate optimal expectile-forecasts at the constant level $\tau = 1/2.85$. For a Gaussian conditional distribution $\mathcal{L}(Y|\mathcal{F}) = \mathcal{N}(\mu, \sigma)$, we obtain the asymmetric forecast

$$x = e_{1/2.85}(\mathcal{N}(\mu, \sigma)) = \mu + \sigma e_{1/2.85}(\mathcal{N}(0, 1)) \approx \mu - \frac{1}{4}\sigma.$$




Figure 2. Size of optimality tests for the optimal full-information forecast in (10). The horizontal line is at the nominal significance level of 0.05.

Let \mathcal{I}_t be the filtration generated by the time series, $\mathcal{I}_t = \sigma(Y_t, Y_{t-1}, \dots)$. A fully informed forecast with the information set $\mathcal{F}_t = \mathcal{I}_{t-1} = \sigma(Y_{t-1}, Y_{t-2}, \dots)$ at time $t-1$, would issue

$$x_t = \frac{1}{2}y_{t-1} - \frac{1}{4}\sigma_t. \quad (10)$$

Applying the standard GMM two-step estimator with the linear specification model using the lagged outcome y_{t-1} as state variable, we perform the overidentifying-restriction tests of forecast optimality from Section 3.3 at significance level 5%. We compare the tests from Section 4 based on linear state-dependent quantiles and expectiles to the flexible spline test introduced in Patton and Timmermann (2007).[†] Using one node only and applying the \mathcal{F}_t -measurable state variable y_{t-1} instead of y_t , the spline test reduces to the expectile linear state-dependent test and the asymptotic results of Section 3 apply. As instruments w_t we use a constant, the forecast, the lagged forecast error, the squared lagged forecast error, and one additional lag of the final three variables.

In Figure 2, we see that the quantile- and expectile-based optimality tests are better calibrated than the spline test. This addresses a known problem of the state-dependent spline test, which “appears to require large samples ($T \geq 1000$) before the test’s size is close to its nominal value, and thus rejections obtained using this test must be interpreted with caution” (Patton and Timmermann, 2007, p. 1183). In contrast to the spline-based estimation, the state-dependent quantile and expectile models provide insightful point estimates and confidence intervals even for moderate sample sizes.

For the power analysis, we construct a 2-step ahead forecast which is optimal with respect to the lagged information set $\mathcal{F}_t = \mathcal{I}_{t-2} = \sigma(Y_{t-2}, Y_{t-3}, \dots)$. For the conditional distribution we obtain $\mu_{Y_t|\mathcal{F}_t} = \mu_{Y_t|\mathcal{I}_{t-2}} = \frac{1}{4}y_{t-2}$ and $\sigma_{Y_t|\mathcal{F}_t}^2 = \sigma_{Y_t|\mathcal{I}_{t-2}}^2 = \frac{23}{20}\sigma_{t-1}^2 + \frac{1}{10}$ as shown in Supplementary Section S3. This produces an optimal expectile-forecast with respect to the information set \mathcal{I}_{t-2} , which is not optimal with respect to the information

[†]One modification has been implemented: The nodes of the splines are located at zero and the average positive and negative forecast errors, instead of the median positive and negative forecast errors, as the medians are close to zero.




Figure 3. Size and power of optimality tests for the 2-step ahead forecast. The solid lines represent size for tests with lagged instruments (nominal level 0.05). The dashed lines represent power for tests with non-lagged instruments.

set \mathcal{I}_{t-1} of variables observable when issuing the forecast. The setting allows us to evaluate the power of the optimality test against information rigidities (Coibion and Gorodnichenko, 2015). A well performing test accepts optimality for lagged instruments based on $\mathcal{F}_t = \mathcal{I}_{t-2}$ according to the 5%-level of the test, and rejects optimality for non-lagged instruments based on \mathcal{I}_{t-1} . The results of this experiment are presented in Figure 3. The quantile- and expectile-based optimality tests are better calibrated and more powerful than the spline-based test, which is strongly oversized for small sample sizes and unable to consistently detect the information rigidity even for large sample sizes.

Hence, our tests are not only better calibrated, but also more powerful. The advantage of the functional-based tests is even greater for more flexible models, where the spline based test would require a large number of instruments. Applying fewer instruments improves the finite sample performance of the quantile and expectile based tests even more (based on additional experiments not shown here). However, for the sake of comparability we have kept the instruments identical across the tests.

5.2. State-dependent quantile-forecasts under different specification models

We generate optimal state-dependent forecasts for the data generating process (9) under the specification models proposed in Section 3.1: a quantile forecast that depends linearly on the current time series value (*linear*: $m(z_t) = \Psi(y_{t-1} - 1)$), one that is subject to periodic deviations (*periodic*: $m(z_t) = \Psi(1 + \sin(\pi t/8))$), and one that is exposed to a break (*break*: $m(z_t) = \Psi(1 - 2 \cdot \mathbb{1}(t \leq T/2))$). To each forecast we apply the overidentifying-restriction tests of forecast optimality from Section 3.3 at level 0.05 based on the three specification models (*linear*, *periodic*, *break*). In the periodic model we consider two parameters and fix the third (period) parameter as we use only three instruments throughout: a constant, the forecast and the lagged outcome. In Table 2, we see the results for forecasts and tests based on quantiles. Additional details and

Table 2. Rejection rates of optimality tests with sample size T based on different specification models at level 0.05.

	true model	hypothesized model		
		linear	periodic	break
$T = 100$	linear	0.06	0.72	0.88
	periodic	0.79	0.08	0.47
	break	0.61	0.30	0.08
$T = 250$	linear	0.06	0.99	1.00
	periodic	1.00	0.07	0.89
	break	0.95	0.64	0.07
$T = 1000$	linear	0.05	1.00	1.00
	periodic	1.00	0.05	1.00
	break	1.00	1.00	0.06

analogous results for tests based on expectiles can be found in Supplementary Section S4. Even for small sample sizes ($T = 100$) the tests are reasonably calibrated and quite powerful with rejection rates between 60% and 90% (with the exception of the periodic model, where the break forecast is rejected in 30% of the cases only). For larger sample sizes ($T \geq 1000$) the optimality tests are almost perfectly calibrated and have full power against the other state-dependent optimal forecasts.

6. Discussion

For point forecasts with unknown directive, we posit that it is preferable to estimate and test the functional quoted by the forecaster, rather than the loss function for reasons of identifiability, interpretability, and ease and efficiency of implementation.

We have introduced state-dependent quantiles and expectiles and have shown that, under optimal forecasts, the functional can be consistently estimated. The asymptotic distributions of the GMM estimator and the overidentifying test statistic can be used to construct flexible tests of forecast optimality and of specific model properties. It is particularly noteworthy that state-dependent functionals allow for the treatment of supposedly misspecified forecasts in a principled manner.

Since Mincer and Zarnowitz (1969) the standard test of optimality is the regression

$$y_t = \beta_0 + \beta_1 x_t + u_t$$

and the joint test that the coefficients in this model are equal to zero and one. Interestingly, this regression model is equivalent to assuming the optimal forecast

$$X = \beta_0 + \beta_1 \mathbb{E}[Y|\mathcal{F}],$$

which yields the identification function $V(x, y) = \beta_0 + \beta_1 x - y$, and applying the GMM-estimator with the instruments $w = (1, x)$.

In a simulation study, we have illustrated that an existing spline-based test is oversized and unlikely to detect information rigidities, while the new estimators yield well calibrated and powerful tests. We further have shown that the GDP forecasts of the Federal Reserve are optimal when viewed as state-dependent quantiles that change with the predicted growth rate.

An important potential application of the new approach is the comparison of a point forecast with unknown directive to other point or probabilistic forecasts. In this situation, the functional represented by the point forecast could be extracted from the probabilistic forecast, and the resulting sets of point forecasts can be compared via any consistent loss function (see Giacomini and White, 2006; Gneiting, 2011).

Acknowledgments

This manuscript has developed from a Diploma thesis written by Patrick Schmidt at Heidelberg University, under the supervision of Matthias Katzfuss and Tilmann Gneiting. We are grateful to Fabian Krüger, Barbara Rossi, the editor, associate editor, and two referees for constructive comments that have greatly improved the manuscript. The work of Patrick Schmidt and Tilmann Gneiting has been partially funded by the Klaus Tschira Foundation and by the European Union Seventh Framework Programme under grant agreement no. 290976. Matthias Katzfuss has been partially supported by US National Science Foundation (NSF) Grant DMS–1521676 and NSF CAREER Grant DMS–1654083.

A. Identification of functionals

Consider the probability space $(\Omega, \mathcal{A}, \mathbb{P})$, where the elements of the sample space Ω are tuples that comprise the realization Y , the point forecast X , and instruments W . We assume that the information set \mathcal{F} is a sub- σ -algebra of \mathcal{A} . If no measure is explicitly mentioned, statements like almost surely refer to \mathbb{P} . For random variables R_1 and R_2 , we simply write $R_1 = R_2$ instead of $R_1 = R_2$ almost surely. In particular, statements like $X = \alpha(Y|\mathcal{F})$ denote \mathbb{P} -almost sure properties.

The terminology on functionals follows Steinwart et al. (2014). Specifically, topological statements on the space of probability distributions \mathcal{P} are with respect to the metric induced by the L_1 -norm.

REGULARITY CONDITIONS A. *There exists a convex set \mathcal{P} of probability measures with bounded Lebesgue densities such that*

- (i) $\mathcal{L}(Y|\mathcal{F}) \in \mathcal{P}$ almost surely, and
- (ii) the functional $\alpha : \mathcal{P} \mapsto \mathbb{R}$ is continuous, locally nonconstant, and has convex level sets.

The conditions A(ii) are met by any functional defined via a continuous, non-trivial loss function. In particular, continuity follows from the Maximum Theorem (Ok, 2007, p. 229), and functionals defined by loss functions have convex level sets (Osband, 1985; Gneiting, 2011).

LEMMA 1 (IDENTIFICATION). *Let X be a forecast for Y based on some information σ -algebra \mathcal{F} , and let $\alpha : \mathcal{P} \mapsto \mathbb{R}$ be a functional. Under Regularity Conditions A, there exists a function V identifying the optimal α -forecast, i.e.,*

$$X = \alpha(Y|\mathcal{F}) \iff \mathbb{E}[V(X, Y)W] = 0 \text{ for all } W \in \mathcal{F}.$$

PROOF. By Theorem 8 in Steinwart et al. (2014), there exists an identification function V , such that $V(t, y)$ exists for $\lambda \times \lambda$ -almost all (t, y) , where λ is the Lebesgue measure, and for all $P \in \mathcal{P}$ it holds that $t = \alpha(P) \iff \mathbb{E}_{Y \sim P}[V(t, Y)] = 0$. In particular,

$$X = \alpha(Y|\mathcal{F}) \iff \mathbb{E}[V(X, Y)|\mathcal{F}] = 0,$$

as $\mathcal{L}(Y|\mathcal{F}) \in \mathcal{P}$ almost surely. As $\mathcal{L}(Y|\mathcal{F})$ is absolutely continuous and X constant given \mathcal{F} , the identification function $V(X, Y)$ exists almost surely.

Any \mathcal{F} -measurable variable W remains constant under conditional integration. Unconditional integration reveals the moment conditions

$$\mathbb{E}[V(X, Y)W] = 0 \quad \text{for all } W \in \mathcal{F}.$$

We show by contradiction that this is also a sufficient condition for $\mathbb{E}[V(X, Y)|\mathcal{F}] = 0$. Assume without loss of generality that $\mathbb{E}[V(X, Y)|\mathcal{F}] > 0$ with positive probability, then the \mathcal{F} -measurable random variable $Z := \mathbb{1}(\mathbb{E}[V(X, Y)|\mathcal{F}] > 0)$ satisfies $\mathbb{E}[V(X, Y)Z] > 0$, contradicting $\mathbb{E}[V(X, Y)W] = 0$ for all $W \in \mathcal{F}$.

B. Consistency of the GMM estimator

We work in the dynamic prediction space setting of Strähl and Ziegel (2017), using the notation introduced in Section 3. The proof covers both the quantile and the expectile parameterization. We define $u = (x, y, z)$ and denote the identification function $V(m(z, \theta), x, y)$ by $V(u, \theta)$, where $V(u, \theta) = \mathbb{1}(y \leq x) - m(z, \theta)$ for quantiles and $V(u, \theta) = (\mathbb{1}(y \leq x) - m(z, \theta))(x - y)$ for expectiles respectively. Then, $g(u, w, \theta) = V(u, \theta)w$ and by Lemma 1, $\mathbb{E}[V(U, \theta_0)|\mathcal{F}] = 0$.

REGULARITY CONDITIONS B.

- (i) $\mathcal{L}(Y|\mathcal{F}) \in \mathcal{P}$ almost surely, where \mathcal{P} is the class of absolutely continuous distributions with strictly positive densities and finite first moments.
- (ii) The state variable Z is \mathcal{F} -measurable.
- (iii) The parameter space $\Theta \subseteq \mathbb{R}^p$ is compact.
- (iv) The specification model $m(z, \theta)$ is continuous on Θ for all z .
- (v) The alternative models differ from the true model, i.e., $m(Z, \theta_0) \neq m(Z, \theta)$ for all $\theta \in \Theta$ with $\theta \neq \theta_0$.
- (vi) The stochastic process $\{U_t | t \in \mathbb{N}\}$ is ergodic (in means) and strictly stationary.†
- (vii) The absolute forecast error has a finite first moment.
- (viii) $\hat{S}_T \rightarrow S$, where S is positive definite.

Condition B(i) ensures that the state-dependent quantiles and expectiles fulfill Regularity Conditions of Lemma 1 for fixed z . Expectiles do not require strictly positive densities. Quantiles do not require finite first moments.

THEOREM 1 (CONSISTENCY). *Let X_t be an optimal state-dependent quantile forecast, i.e., $X_t = q_{m(Z_t, \theta_0)}(Y_t|\mathcal{F}_t)$ for $t = 1, 2, \dots$. Under Regularity Conditions B, there exists an \mathcal{F}_t -measurable instrumental variable W_t such that the GMM estimator defined in (6) is consistent for the true parameter $\theta_0 \in \Theta$:*

$$\hat{\theta}_T \rightarrow \theta_0, \quad \text{as } T \rightarrow \infty.$$

Analogously, the expectile based estimator is consistent if the point forecast is a state-dependent expectile, $X_t = e_{m(Z_t, \theta_0)}(Y_t|\mathcal{F}_t)$.

†Strict stationarity means that the distribution of $(U_t, U_{t+1}, \dots, U_{t+s})$ does not depend on t for any s , and (mean) ergodicity implies that $\frac{1}{T} \sum_{t=1}^T a(U_t)$ converges in probability to $\mathbb{E}[a(U_t)]$ for measurable functions $a(\cdot)$ with $\mathbb{E}[|a(U_t)|] < \infty$ (Newey and McFadden, 1986).

LEMMA 2 (UNIQUENESS). *Given B(i), B(ii) and B(v), it holds for all $t = 1, 2, \dots$ that*

$$\mathbb{E}[V(U_t, \theta)|\mathcal{F}_t] = 0 \iff \theta = \theta_0.$$

PROOF. We first consider quantiles. Under B(i) the state-dependent quantile is well-defined and fulfills Regularity Conditions A for fixed z . By Lemma 1, it follows that $\mathbb{E}[V(U, \theta_0)|\mathcal{F}] = 0$. For any $\theta \in \Theta$ with $\theta \neq \theta_0$ it holds that

$$\begin{aligned} \mathbb{E}[V(U, \theta)|\mathcal{F}] &= \mathbb{E}[m(Z, \theta) - \mathbb{1}(X > Y)|\mathcal{F}] \\ &= m(Z, \theta) - \mathbb{E}[\mathbb{1}(q_{m(Z, \theta_0)}(Y|\mathcal{F}) > Y)|\mathcal{F}] \\ &= m(Z, \theta) - m(Z, \theta_0) \neq 0, \end{aligned}$$

because competing models differ almost surely and $Z \in \mathcal{F}$ by B(ii) and B(v).

Analogously, expectiles are well-defined under B(i) for fixed z and by Lemma 1 it follows that $\mathbb{E}[V(U, \theta_0)|\mathcal{F}] = 0$. The moment condition is

$$\mathbb{E}[\mathbb{1}(X > Y)(X - Y) - m(Z, \theta)(X - Y)(1 - 2 \cdot \mathbb{1}(X < Y))|\mathcal{F}] = 0.$$

Under B(ii) and defining $P = Y|\mathcal{F}$, we separate the integral at point $Y = x$, to obtain

$$\int_{-\infty}^x (x - y)dP(y) - m(z, \theta) \left(\int_{-\infty}^x (x - y)dP(y) - \int_x^{\infty} (x - y)dP(y) \right) = 0,$$

which is equivalent to

$$(1 - m(z, \theta)) \int_{-\infty}^x (x - y)dP(y) = m(z, \theta) \int_x^{\infty} (y - x)dP(y). \quad (11)$$

The $m(z, \theta)$ -expectile for the distribution P is the unique solution of (11) (Newey and Powell, 1987). If $\theta \neq \theta_0$, it follows by B(v) that $m(Z, \theta) \neq m(Z, \theta_0)$, which implies that $X \neq e_{m(Z, \theta)}(Y|\mathcal{F})$. Consequently, moment condition (11) does not hold.

PROOF (THEOREM 1). We verify that there exist instruments W such that the conditions of Theorem 2.6 of Newey and McFadden (1986, pp. 2132–2133) are satisfied, which directly implies consistency.

From Lemma 2, it follows that $f(\theta) := \mathbb{E}[V(U, \theta)|\mathcal{F}]$ is the constant zero function if and only if $\theta = \theta_0$. Let $\theta \neq \theta_0$ and $A := \mathbb{1}(m(Z, \theta_0) > m(Z, \theta))$. For fixed θ , it holds that $A \in \mathcal{F}$, because by B(ii) Z is \mathcal{F} -measurable. The same holds true for its complement A^c , as \mathcal{F} is a σ -algebra. By B(v), $m(Z, \theta_0) \neq m(Z, \theta)$ for all $\theta \in \Theta$ with $\theta \neq \theta_0$. It follows that either A or A^c have positive weight. We define $W := (\mathbb{1}(A), -\mathbb{1}(A^c))$. As $m(Z, \theta_0) > m(Z, \theta) \iff X > q_{m(Z, \theta)}(Y|\mathcal{F}) \iff \mathbb{E}[V(U, \theta)|\mathcal{F}] > 0$, it follows that $f(\theta) > 0 \iff \mathbb{1}(A) = 1$. We have found a random vector $W \in \mathcal{F}$ such that for $\theta \neq \theta_0$

$$\mathbb{E}[g(U, \theta)] = \mathbb{E}[\mathbb{E}[V(U, \theta)W|\mathcal{F}]] = \mathbb{E}[f(\theta)W] \neq 0,$$

as $f(\theta)W > 0$ elementwise by definition. Consequently, there exists an \mathcal{F} -measurable instrumental vector W such that the unique identification property i) of Newey and McFadden (1986, Thm. 2.6) holds. By B(vi), W is ergodic and stationary, as a measurable function of the ergodic and stationary variable Z .

As S is positive definite by B(viii), its inverse S^{-1} exists and is also positive definite. It follows directly that, $S^{-1}\mathbb{E}[g(U, \theta)] = 0 \iff \mathbb{E}[g(U, \theta)] = 0$.

As the parameter space Θ is compact by B(iii), and the applied functional is in our parameterized class, it follows that $\theta_0 \in \Theta$ and ii) of Newey and McFadden (1986, Thm. 2.6) is satisfied.

The continuity of $g(u, w, \theta)$ in θ (Newey and McFadden, 1986, Thm. 2.6 iii)) follows directly from B(iv): For quantiles, it holds that $g(u, w, \theta) = (\mathbb{1}(x > y) - m(z, \theta))w$, hence g is continuous in the parameter θ , as a composition of continuous functions. For expectiles, it holds that $g(u, w, \theta) = [\mathbb{1}(x > y) - m(z, \theta)(1 - 2 \cdot \mathbb{1}(x < y))](x - y)w$, which again is a continuous function in θ .

Finally, $\sup_{\theta \in \Theta} \|g(u, w, \theta)\| \leq \|w\| < \infty$ for quantiles and $\sup_{\theta \in \Theta} \|g(u, w, \theta)\| \leq \|w\| \|(x - y)\| \leq |x - y|$ for expectiles, which implies $\mathbb{E}[\sup_{\theta \in \Theta} \|g(U, W, \theta)\|] < \infty$ by B(vii). Thus, Condition iv) of Newey and McFadden (1986, Thm. 2.6) is satisfied, and the GMM estimator is consistent.

References

- Capistrán, C. (2008) Bias in Federal Reserve inflation forecasts: Is the Federal Reserve irrational or just cautious? *Journal of Monetary Economics*, **55**, 1415–1427.
- Caunedo, J., DiCecio, R., Komunjer, I. and Owyang, M. T. (2013) Asymmetry and Federal Reserve forecasts. *Federal Reserve Bank of St. Louis Working Paper Series*.
- Christodoulakis, G. A. and Mamatzakis, E. C. (2008) An assessment of the EU growth forecasts under asymmetric preferences. *Journal of Forecasting*, **27**, 483–492.
- Coibion, O. and Gorodnichenko, Y. (2015) Information rigidity and the expectations formation process: a simple framework and new facts. *American Economic Review*, **105**, 2644–2678.
- Diebold, F. X. and Lopez, J. A. (1996) Forecast evaluation and combination. In *Statistical Methods in Finance* (eds. G. S. Maddala and C. R. Rao), vol. 14 of *Handbook of Statistics*, 241–268. Elsevier.
- Ehm, W., Gneiting, T., Jordan, A. and Krüger, F. (2016) Of quantiles and expectiles: consistent scoring functions, Choquet representations and forecast rankings. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **78**, 505–562.
- Elliott, G., Komunjer, I. and Timmermann, A. (2005) Estimation and testing of forecast rationality under flexible loss. *The Review of Economic Studies*, **72**, 1107–1125.
- (2008) Biases in macroeconomic forecasts: irrationality or asymmetric loss? *Journal of the European Economic Association*, **6**, 122–157.
- Elliott, G. and Timmermann, A. (2008) Economic forecasting. *Journal of Economic Literature*, **46**, 3–56.
- Engelberg, J., Manski, C. F. and Williams, J. (2009) Comparing the point predictions and subjective probability distributions of professional forecasters. *Journal of Business & Economic Statistics*, **27**, 30–41.
- Fissler, T. and Ziegel, J. F. (2016) Higher order elicibility and Osband’s principle. *The Annals of Statistics*, **44**, 1680–1707.
- Fritsche, U., Pierdzioch, C., Rülke, J.-C. and Stadtmann, G. (2015) Forecasting the Brazilian Real and the Mexican Peso: Asymmetric loss, forecast rationality, and forecaster herding. *International Journal of Forecasting*, **31**, 130–139.
- Giacomini, R. and White, H. (2006) Tests of conditional predictive ability. *Econometrica*, **74**, 1545–1578.

- Gneiting, T. (2011) Making and evaluating point forecasts. *Journal of the American Statistical Association*, **106**, 746–762.
- Gneiting, T., Balabdaoui, F. and Raftery, A. E. (2007) Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **69**, 243–268.
- Gneiting, T. and Katzfuss, M. (2014) Probabilistic forecasting. *Annual Review of Statistics and Its Application*, **1**, 125–151.
- Gneiting, T. and Ranjan, R. (2013) Combining predictive distributions. *Electronic Journal of Statistics*, **7**, 1747–1782.
- Greene, W. H. (2012) *Econometric Analysis*. Upper Saddle River, NJ: Prentice Hall, 7th edn.
- Guler, K., Ng, P. T. and Xiao, Z. (2017) Mincer-Zarnowitz quantile and expectile regressions for forecast evaluations under asymmetric loss functions. *Journal of Forecasting*, **36**, 651–679.
- Hansen, L. P. (1982) Large sample properties of generalized method of moments estimators. *Econometrica*, **50**, 1029–1054.
- Heinrich, C. (2014) The mode functional is not elicitable. *Biometrika*, **101**, 245–251.
- Horowitz, J. L. and Manski, C. F. (2006) Identification and estimation of statistical functionals using incomplete data. *Journal of Econometrics*, **132**, 445–459.
- Huber, P. J. and Ronchetti, E. M. (2009) *Robust Statistics*. New Jersey: John Wiley & Sons, Inc., 2nd edn.
- Komunjer, I. and Owyang, M. T. (2012) Multivariate forecast evaluation and rationality testing. *Review of Economics and Statistics*, **94**, 1066–1080.
- Körding, K. P. and Wolpert, D. M. (2004) The loss function of sensorimotor learning. *Proceedings of the National Academy of Sciences of the United States of America*, **101**, 9839–9842.
- Krol, R. (2013) Evaluating state revenue forecasting under a flexible loss function. *International Journal of Forecasting*, **29**, 282–289.
- Lieli, R. P. and Stinchcombe, M. B. (2013) On the recoverability of forecasters’ preferences. *Econometric Theory*, **29**, 517–544.
- Manski, C. F. (2016) Interpreting point predictions: Some logical issues. *Foundations and Trends in Accounting*, **10**, 238–261.
- Mincer, J. A. and Zarnowitz, V. (1969) The evaluation of economic forecasts. In *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance* (ed. J. A. Mincer), 3–46. National Bureau of Economic Research.
- Newey, W. K. and McFadden, D. (1986) Large sample estimation and hypothesis testing. In *Handbook of Econometrics* (eds. R. F. Engle and D. McFadden), vol. 4 of *Handbook of Econometrics*, chap. 36, 2111–2245. Elsevier.
- Newey, W. K. and Powell, J. L. (1987) Asymmetric least squares estimation and testing. *Econometrica*, **55**, 819–847.

- Newey, W. K. and West, K. D. (1987) A simple, positive semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, **55**, 703–708.
- Ok, E. A. (2007) *Real Analysis with Economic Applications*. Princeton University Press.
- Osband, K. H. (1985) *Providing Incentives for Better Cost Forecasting*. Ph.D. thesis, University of California, Berkely.
- Patton, A. J. and Timmermann, A. (2007) Testing forecast optimality under unknown loss. *Journal of the American Statistical Association*, **102**, 1172–1184.
- Pierdzioch, C., Rülke, J.-C. and Stadtmann, G. (2013) Oil price forecasting under asymmetric loss. *Applied Economics*, **45**, 2371–2379.
- Savage, L. J. (1971) Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, **66**, 783–801.
- Sims, C. R. (2015) The cost of misremembering: inferring the loss function in visual working memory. *Journal of Vision*, **15**, 2–2.
- Steinwart, I., Pasin, C., Williamson, R. and Zhang, S. (2014) Elicitation and identification of properties. In *Proceedings of The 27th Conference on Learning Theory* (eds. M. F. Balcan, V. Feldman and C. Szepesvári), vol. 35 of *Proceedings of Machine Learning Research*, 482–526.
- Strähl, C. and Ziegel, J. (2017) Cross-calibration of probabilistic forecasts. *Electronic Journal of Statistics*, **11**, 608–639.
- Wang, Y. and Lee, T.-H. (2014) Asymmetric loss in the Greenbook and the Survey of Professional Forecasters. *International Journal of Forecasting*, **30**, 235–245.