



***The World's Largest Open Access Agricultural & Applied Economics Digital Library***

**This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.**

**Help ensure our sustainability.**

Give to AgEcon Search

AgEcon Search  
<http://ageconsearch.umn.edu>  
[aesearch@umn.edu](mailto:aesearch@umn.edu)

*Papers downloaded from AgEcon Search may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

*No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.*

# Conditional evaluation of predictive models: The `cspa` command

Jia Li	Zhipeng Liao
Singapore Management University	University of California–Los Angeles
Singapore	Los Angeles, CA
jiali@smu.edu.sg	zhipeng.liao@ucla.edu
Rogier Quaedvlieg	Wenyu Zhou
European Central Bank	Zhejiang University
Frankfurt, Germany	Hangzhou, China
rogier.quaedvlieg@ecb.europa.eu	wenyuzhou@intl.zju.edu.cn

**Abstract.** In this article, we introduce a new command, `cspa`, that implements the conditional superior predictive ability test developed in Li, Liao, and Quaedvlieg (2022, *Review of Economic Studies* 89: 843–875). With the conditional performance of predictive methods measured nonparametrically by the conditional expectation functions of their predictive losses, we test the null hypothesis that a benchmark model weakly outperforms a collection of competitors uniformly across the conditioning space. The proposed command can implement this test for both independent cross-sectional data and serially dependent time-series data. Confidence sets for the most superior model can be obtained by inverting the test, for which the `cspa` command also offers a convenient implementation.

**Keywords:** st0696, `cspa`, conditional moment inequality, forecast evaluation, functional inference, series estimation

## 1 Introduction

Making accurate predictions is a central task of data analysis. Economists in central banks routinely update their forecasts on a broad range of macroeconomic indicators to inform policymaking. Investors in financial markets devote a tremendous amount of effort to predict asset prices' movements. Data analytics for various socioeconomic activities have also stimulated new generations of predictive models on the “micro” level for firms, communities, households, and individuals. At the same time, numerous machine-learning techniques have been developed to address complex forecasting problems in the era of “big data”.

In view of the ever-growing abundance of potentially good predictive methods, evaluating their relative performance is evidently of great practical importance. The most intuitive approach, which is also commonly adopted in machine-learning studies, is to directly compare the forecasting accuracy measured by certain loss functions. Such evaluation may be further formalized as a statistical hypothesis testing problem. Arguably the most popular test in economic applications is the pseudo out-of-sample Diebold–Mariano test (Diebold and Mariano 1995) for the null hypothesis that the ex post losses

(for example, the squared prediction error) of two predictive methods are equal in expectation. Diebold and Mariano's proposal is effectively a *t* test that may be implemented by regressing the observed loss differential between the competing methods on a constant term and checking whether the estimated intercept is significantly different from zero. This can easily be done in Stata via `regress` (see [R] `regress`) for independent data or, more generally, `newey` (see [TS] `newey`) for data with serial dependence.

Although the Diebold–Mariano test provides a simple way of forecasting evaluation, a more practically relevant question would be to investigate whether a specific method generates more accurate forecasts than its competitors. Following this direction, the test for unconditional superior predictive ability was first developed by White (2000) and later refined by Hansen (2005). The null hypothesis of such a test states that a benchmark method (weakly) outperforms a set of competing alternatives.

Because the aforementioned tests are based on the unconditional average performance of the competing predictions, they invariably “integrates out” the heterogeneity across subsample periods (for example, expansion or recession episodes) or subpopulations (for example, income or age groups) and hence may be too coarse for certain practical applications. Unraveling the state-dependent or characteristic-dependent relative performance among predictive models naturally calls for a conditional evaluation approach. The first attempt in this aspect, to our best knowledge, traces back to the conditional equal predictive ability (CEPA) test proposed by Giacomini and White (2006). Under the CEPA null hypothesis, the conditional expected loss functions of different forecasting methods are identically the same across all conditioning states.

In this article, we propose a new command, `cspa`, that offers a convenient implementation for the conditional superior predictive ability (CSPA) test recently proposed by Li, Liao, and Quaedvlieg (2022). The null hypothesis of interest states that the conditional expected loss of a benchmark predictive model is no larger than those of the competing models uniformly across all conditioning states (specified by a conditioning variable chosen *a priori* by the user). In this sense, the CSPA null hypothesis asserts that the benchmark is a uniformly weakly dominating method among all predictive models under consideration. “Passing” the test indicates that the benchmark method is likely to perform well not only on average but also across all subpopulations “sliced” by the conditioning variable chosen by the user, which may be a macroeconomic cyclical indicator, a financial risk measure, or an individual characteristic, depending on the empirical context.

The main functionality of the proposed `cspa` command is to implement the CSPA test for a given benchmark against a collection of competitors. A rejection indicates that some competitor outperforms the benchmark over some conditioning states. A nonrejection, on the other hand, suggests that the benchmark is weakly dominating. Moreover, by rotating the benchmark role across all models, we may form the confidence set for the most superior (CSMS) as the collection of all nonrejected benchmarks. This operation can also be accomplished easily by calling the `csms` option. We shall demonstrate how to carry out these inferential tasks in an empirical example on asset volatility forecasting using a dataset from Li, Liao, and Quaedvlieg (2022).

The remainder of this article is organized as follows. Section 2 provides a brief review on the CSPA testing procedure and the underlying intuition. Section 3 documents the syntax of the `cspa` command and available options. Section 4 demonstrates the key functionalities of the proposed command in an empirical example. Section 5 concludes.

## 2 The conditional evaluation method

This section provides a brief review of the CSPA test developed in Li, Liao, and Quaedvlieg (2022). Section 2.1 describes the setting. Section 2.2 details the CSPA testing procedure. To simplify the discussion, we mainly focus on the time-series setting, with the understanding that random samples with independent observations may be considered as a special “time series” with no serial dependence.

### 2.1 The setting and hypotheses of interest

Let  $(F_t^\dagger)_{1 \leq t \leq n}$  denote the time series to be predicted. Correspondingly, we consider a series of benchmark predictions,  $(F_{0,t})_{1 \leq t \leq n}$ , and a collection of  $J$  competing alternatives,  $(F_{j,t})_{1 \leq t \leq n}$ ,  $1 \leq j \leq J$ . The performance of these predictive methods is gauged by a user-specified loss function,  $L(\cdot, \cdot)$ , so the loss for model  $j$  at period  $t$  is  $L(F_t^\dagger, F_{j,t})$ . Commonly used loss functions include the quadratic loss, absolute deviation loss, and Stein’s loss.<sup>1</sup> The performance of the benchmark method relative to the  $j$ th competing alternative is thus measured by the loss differential series

$$Y_{j,t} = L(F_t^\dagger, F_{j,t}) - L(F_t^\dagger, F_{0,t})$$

Its conditional mean function, given a user-specified conditioning variable  $X_t$ , is further defined as

$$h_j(x) \equiv \mathbb{E}(Y_{j,t} | X_t = x)$$

Note that  $h_j(x) \geq 0$  indicates that the benchmark method is expected to (weakly) outperform the  $j$ th competitor conditional on  $X_t = x$ .

The null hypothesis of the CSPA test asserts that

$$H_0 : h_j(x) \geq 0, \quad \text{for all } x \in \mathcal{X}, \quad 1 \leq j \leq J \tag{1}$$

where  $\mathcal{X}$  is the support of  $X_t$ . Under the null hypothesis, the benchmark method outperforms all competing alternatives uniformly across all conditioning states. Evidently, by the law of iterated expectations, this also implies that the unconditional expected loss of the benchmark is lower than those of the competing methods [that is,  $\mathbb{E}(Y_{j,t}) \geq 0$ ]. However, the CSPA null hypothesis is generally much more stringent than its unconditional counterpart. As such, the (uniform) conditional dominance criterion may help the researcher discriminate competing predictive methods that appear “unconditionally similar”, and thus it complements conventional evaluation methods such as the Diebold–Mariano test.

---

1. Stein’s loss function is defined in (3) on page 931.

## 2.2 The testing procedure

We now detail how the `cspa` command implements the CSPA test in the background. The first step is to nonparametrically estimate the conditional expectation functions of the loss differentials [that is,  $\{h_j(\cdot) : 1 \leq j \leq J\}$ ] by running series regressions. Specifically, let  $\mathbf{p}(x) = (p_1(x), \dots, p_m(x))^\top$  be an  $m$ -dimensional vector of approximating basis functions. For each  $j$ , we regress  $Y_{j,t}$  on  $\mathbf{p}(X_t)$  with the resulting regression coefficient given by

$$\hat{\mathbf{b}}_j = \hat{\mathbf{Q}}^{-1} \left\{ n^{-1} \sum_{t=1}^n \mathbf{p}(X_t) Y_{j,t} \right\} \quad \text{where} \quad \hat{\mathbf{Q}} = n^{-1} \sum_{t=1}^n \mathbf{p}(X_t) \mathbf{p}(X_t)^\top$$

The nonparametric series estimator for  $h_j(\cdot)$  is then constructed as

$$\hat{h}_j(\cdot) = \mathbf{p}(\cdot)^\top \hat{\mathbf{b}}_j$$

The underlying nonparametric inference theory requires the number of series terms  $m \rightarrow \infty$  in large samples, so the unknown conditional expectation functions can be well approximated.

The current version of `cspa` uses Legendre polynomials to form the approximating functions  $\mathbf{p}(X_t)$ . An important property of the Legendre polynomials is that they are orthogonal on the  $[-1, 1]$  interval with respect to the uniform distribution. This orthogonality property helps mitigate the multicollinearity among series terms and hence improves the numerical stability of the estimation procedure in finite samples. Other types of orthogonal series basis may also be adopted to serve the same purpose, which is left for future development. To better exploit the orthogonality of Legendre polynomials, we recommend performing a preliminary transformation on the conditioning variable  $X_t$  to make it approximately uniformly distributed on the  $[-1, 1]$  interval; see the `method()` option for available choices provided in `cspa`.

To carry out the test, one must account for the sampling variability of the functional estimates  $\hat{h}_j(\cdot)$ . Let  $\hat{u}_{j,t} = Y_{j,t} - \hat{h}_j(X_t)$  be the residual from the  $j$ th regression, and collect them in a  $J$ -dimensional column vector  $\hat{\mathbf{u}}_t$ . We estimate the joint variance-covariance matrix for the series regression coefficients  $\hat{\mathbf{b}}_j$ ,  $1 \leq j \leq J$ , using

$$\hat{\Omega} \equiv (\mathbf{I}_J \otimes \hat{\mathbf{Q}})^{-1} \hat{\mathbf{A}} (\mathbf{I}_J \otimes \hat{\mathbf{Q}})^{-1}$$

where  $\mathbf{I}_J$  denotes the  $J \times J$  identity matrix,  $\otimes$  is the Kronecker product, and  $\hat{\mathbf{A}}$  is a Newey-West estimator for the “long-run” covariance matrix of the  $Jm$ -dimensional vector  $\hat{\mathbf{u}}_t \otimes \mathbf{p}(X_t)$ . The lag parameter for the Newey-West estimator may be set via the `lag()` option in `cspa`; also see [TS] `newey`. Note that if it is known a priori that there is no serial correlation in the data (for example, when the observations are assumed to be independent), one may set `lag(0)`, which is also the default option in `cspa`. We further partition  $\hat{\Omega}$  into  $J \times J$  blocks of  $m \times m$  submatrices. The standard error function of  $\hat{h}_j(\cdot)$  is then estimated as

$$\widehat{\sigma}_j(x) \equiv \left\{ \mathbf{p}(x)^\top \widehat{\Omega}(j,j) \mathbf{p}(x) \right\}^{1/2}$$

where  $\widehat{\Omega}(j,j)$  is the  $(j,j)$  block extracted from the aforementioned partition of  $\widehat{\Omega}$ .

At a significance level  $\alpha$ , the rejection decision of the CSPA test is determined as follows.

**Step 1.** Simulate a  $Jm$ -dimensional normal random vector  $(\xi_1^{*\top}, \dots, \xi_J^{*\top})^\top \sim \mathcal{N}(\mathbf{0}, \widehat{\Omega})$ , where each  $\xi_j^*$  is  $m$  dimensional. Set  $\widehat{t}_j^*(x) \equiv \mathbf{p}(x)^\top \xi_j^* / \widehat{\sigma}_j(x)$ .

**Step 2.** Repeat step 1 many times. For some constant  $z > 0$ , define  $\widehat{K}$  as the  $1 - z / \log(n)$  quantile of  $\max_{1 \leq j \leq J} \sup_{x \in \mathcal{X}} \widehat{t}_j^*(x)$  in the simulated sample and then set

$$\begin{aligned} \widehat{V} \equiv & \left\{ (j, x) : \widehat{h}_j(x) \leq \min_{1 \leq j \leq J} \inf_{x \in \mathcal{X}} \left\{ \widehat{h}_j(x) + n^{-1/2} \widehat{K} \widehat{\sigma}_j(x) \right\} \right. \\ & \left. + 2n^{-1/2} \widehat{K} \widehat{\sigma}_j(x) \right\} \end{aligned}$$

The default value of  $z$  is 0.1, which may be modified by calling the `ais()` option.

**Step 3.** Set  $\widehat{k}_{1-\alpha}$  as the  $(1 - \alpha)$  quantile of  $\sup_{(j,x) \in \widehat{V}} \widehat{t}_j^*(x)$ . Reject the null hypothesis (1) if and only if

$$\widehat{\eta}_{1-\alpha} \equiv \min_{1 \leq j \leq J} \inf_{x \in \mathcal{X}} \left\{ \widehat{h}_j(x) + n^{-1/2} \widehat{k}_{1-\alpha} \widehat{\sigma}_j(x) \right\} < 0 \quad (2)$$

Let us clarify some intuition underlying this testing procedure. The set  $\widehat{V}$  described in step 2 implements an adaptive inequality selection (corresponding with the `ais()` option). It can be shown that, with probability approaching 1,  $\widehat{V}$  contains all  $(j, x)$ 's that minimize  $h_j(x)$ . From (1), we can easily see that whether the null hypothesis holds is solely determined by the functions' values at these extremum points. The selection step focuses the test on the relevant conditioning region and so improves its power. Step 3 then investigates the validity of the null hypothesis by examining the uniform upper confidence bounds for  $h_j(\cdot)$ . We note that the rejection decision described in (2) naturally admits a graphical representation. Indeed, the functional estimates  $\widehat{h}_j(\cdot) + n^{-1/2} \widehat{k}_{1-\alpha} \widehat{\sigma}_j(\cdot)$  are  $1 - \alpha$  upper confidence bounds for  $h_j(\cdot)$  uniformly over the selected set  $\widehat{V}$ . If some of these upper bounds dip below zero over some part of the conditioning space, we interpret it as significant statistical evidence against the null hypothesis [that is,  $h_j(x) \geq 0$  for all  $(j, x)$ ], which leads to a formal rejection.

For some applications, it may be more natural to treat all competing predictive methods “symmetrically”, rather than picking a particular one as the benchmark. In this situation, one may naturally rotate the benchmark role across all competing methods and then collect all nonrejected benchmarks in a set,

$$\begin{aligned} \widehat{M}_{1-\alpha} = & \{0 \leq j \leq J : \text{the } \alpha\text{-level CSPA test} \\ & \text{with method } j \text{ as the benchmark does not reject}\} \end{aligned}$$

This operation is formally an inversion of the CSPA test, and, by the duality between tests and confidence sets,  $\widehat{M}_{1-\alpha}$  is a  $1-\alpha$  level CSMS (which collects methods with the lowest conditional expected loss uniformly across all conditioning states). The implementation can be easily carried out via `cspa` by calling the `csms` option.

## 3 The `cspa` command

This section documents the syntax and functionalities of the `cspa` command. The command requires the `moremata` (Jann 2005) package, which may be installed in the command line via `ssc install moremata`.

### 3.1 Syntax

The syntax of the `cspa` command is

```
cspa condvar benchmark competitors [if] [in] [, ais(#) lag(#) m(#)
method(transtype) siglevel(#) ngrid(#) triml(#) trimr(#) mc(#) csms
plot plotu detail(color) excel]
```

*condvar* is the conditioning variable; *benchmark* is the loss of the benchmark; and *competitors* is a list of loss variables associated with the competitors.

### 3.2 Options

`ais(#)` specifies the degree of adaptive inequality selection. The default is `ais(0.1)`. The selection may be disabled by setting `ais(0)`.

`lag(#)` specifies the number of lags for computing the Newey–West estimator  $\widehat{A}_n$ . The default is `lag(0)`.

`m(#)` specifies the number of Legendre polynomial terms used in series estimation. The default is given by the integer part of  $\max(4, n^{1/5})$ .

`method(transtype)` specifies the transformation implemented on the conditioning variable. The main purpose of doing so is to make the regressors approximately orthogonal, which generally improves the numerical stability of the series regression, especially when many series terms are included. The approximating functions are Legendre polynomials of the transformed variable. The current version supports the following transformation methods, with `method(rank)` being the default.

`none`: no transformation

`affine`: affine transformation  $x \mapsto 2\{x - \min(x)\} / \{\max(x) - \min(x)\} - 1$

**normal**: normal transformation  $x \mapsto 2\Phi\{(x - \bar{x})/\sigma\} - 1$ , where  $\bar{x}$  and  $\sigma$  are the sample mean and standard deviation of  $x$  and  $\Phi$  is the cumulative distribution function of the standard normal distribution

**lognormal**: log-normal transformation  $x \mapsto 2\Phi\{(\log x - \bar{\log x})/\Sigma\} - 1$ , where  $\bar{\log x}$  and  $\Sigma$  are the sample mean and standard deviation of  $\log x$  and  $\Phi$  is the cumulative distribution function of the standard normal distribution

**rank**:  $x \mapsto 2q(x) - 1$ , where  $q(x)$  is the empirical quantile of  $x$

**siglevel(#)** specifies the significance level (in percent) for the CSPA test. The default is **siglevel(5)**.

**ngrid(#)** specifies the number of grid points used for discretizing the support of the conditioning state variable. The default is **ngrid(1000)**.

**triml(#)** sets the left limit of the conditioning region  $\mathcal{X}$  to be the  $\#$  empirical quantile of *condvar*. The default is **triml(0)**.

**trimr(#)** sets the right limit of the conditioning region  $\mathcal{X}$  to be the  $1 - \#$  empirical quantile of *condvar*. The default is **trimr(0)**.

**mc(#)** specifies the number of Monte Carlo simulations used to construct the critical values. The default is **mc(5000)**.

**csms** calculates the CSMS by rotating the benchmark role across all competing methods.

**plot** produces plots of the lower envelope of estimated conditional mean functions of loss differentials and their upper confidence bounds, with the transformed conditioning variable plotted on the horizontal axis.

**plotu** produces plots of the lower envelope of estimated conditional mean functions of loss differentials and their upper confidence bounds, with the original conditioning variable plotted on the horizontal axis.

**detail(*color*)** adds the estimated conditional mean functions of loss differentials onto the plot generated by **plot** or **plotu**. The color of these functions is specified by *color*.

**excel** creates an Excel file that contains the requisite estimates for producing the plots generated by **plot** or **plotu**.

### 3.3 Stored results

`cspa` stores the following in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(ts)</code>	CSPA test statistic
<code>e(pvalue)</code>	<i>p</i> -value of the CSPA test

Macros

<code>e(cmd)</code>	<code>cspa</code>
<code>e(condvar)</code>	name of the conditioning variable
<code>e(loss)</code>	name of the loss variables
<code>e(method)</code>	transformation method

Matrices

<code>e(xgrid)</code>	grid points of the conditioning variable
<code>e(h_hat)</code>	estimates of the conditional expected loss differentials
<code>e(lowerenvelope)</code>	estimate of the lower envelope of the conditional expected loss differentials
<code>e(cb)</code>	estimate of the confidence bound

## 4 An empirical example

### 4.1 Data description

Our real-data demonstration is based on a dataset from the empirical study of Li, Liao, and Quaedvlieg (2022). `rv.dta` contains a time series of daily realized volatilities of the Boeing Company, computed as the sum of squared five-minute returns within regular trading hours, and six time series of one-day-ahead volatility forecasts that are generated using a first-order autoregressive process, AR(1), an AR(22), and an AR(22) with lasso selection, the heterogeneous autoregressive (HAR) model of Corsi (2009), the heterogeneous autoregressive quarticity (HARQ) model developed in Bollerslev, Patton, and Quaedvlieg (2016), and an autoregressive fractionally integrated moving average (ARFIMA) model. These forecasts are formed under a rolling-window scheme with 1,000 daily observations, and the evaluation sample ranges from May 2001 to December 2013, resulting in 3,180 daily forecasts. We refer the reader to Li, Liao, and Quaedvlieg (2022) for details on their constructions. The dataset also contains a time series of lagged Chicago Board Options Exchange volatility indexes (VIX) that shall be used as the conditioning variable  $X_t$  for conducting the CSPA test.

Following Li, Liao, and Quaedvlieg (2022), we use Stein's loss function to measure the performance of the forecasts; that is,

$$L(F_t^\dagger, F_{j,t}) = \frac{F_{j,t}}{F_t^\dagger} - \log\left(\frac{F_{j,t}}{F_t^\dagger}\right) - 1 \quad (3)$$

This is a natural choice because volatility is a scale parameter and Stein's loss is commonly used in the study of scale problems. The losses of the six volatility forecasts are then generated as follows:

```

. *** Load data ***
. use rv

. *** Calculate loss ***
. generate ar1      = (rv_ar1/true)      - log((rv_ar1/true)) - 1
. generate ar22     = (rv_ar22/true)     - log((rv_ar22/true)) - 1
. generate ar22lasso = (rv_ar22lasso/true) - log((rv_ar22lasso/true)) - 1
. generate har      = (rv_har/true)      - log((rv_har/true)) - 1
. generate harq     = (rv_harq/true)     - log((rv_harq/true)) - 1
. generate arfima   = (rv_arfima/true)   - log((rv_arfima/true)) - 1

```

## 4.2 One-versus-one CSPA test

We illustrate the most basic usage of the *cspa* command by conducting a one-versus-one conditional evaluation for two competing models: AR(1) and HAR. Specifically, we perform the CSPA test with one of them as the benchmark and the other as the competitor (so  $J = 1$ ). The tests are implemented as follows:

```
. cspa vix ar1 har, lag(11) plot
```

Transformation on vix: Rank.

Benchmark	CSPA Test (5%)	p-value
ar1	-0.1029 (reject)	0.0002

```
. cspa vix har ar1, lag(11) plot
```

Transformation on vix: Rank.

Benchmark	CSPA Test (5%)	p-value
har	0.0405 (non-reject)	0.9998

The output table of *cspa* reports the test statistic  $\hat{\eta}_{1-\alpha}$  at the user-specified significance level  $\alpha$ , which by default is 5%. The null hypothesis that the benchmark uniformly weakly dominates the competitor is rejected when  $\hat{\eta}_{1-\alpha} < 0$ . The table also reports the *p*-value of the test. Looking at the tables above, we see that the CSPA null hypothesis with AR(1) being the benchmark is strongly rejected, but the test does not reject the CSPA null for HAR at any conventional significance level, suggesting that HAR is the superior predictive method.

With *plot* specified, *cspa* also plots the estimated conditional expected loss differential function  $\hat{h}_1(\cdot)$  along with the upper confidence bound  $\hat{h}_1(\cdot) + n^{-1/2}\hat{k}_{1-\alpha}\hat{\sigma}_1(\cdot)$  as shown in figure 1. Recall that a negative loss differential indicates that the benchmark underperforms the competitor, and vice versa. Moreover, the CSPA test rejects the null hypothesis if the upper confidence bound dips below zero over some part of the conditioning space. From the left panel of figure 1, we see that AR(1) actually underperforms HAR over the entire conditioning space with high statistical significance. Meanwhile, the upper confidence bound plotted on the right panel is always above zero, which is

consistent with the fact that the CSPA test does not reject the null hypothesis that HAR uniformly weakly dominates AR(1).

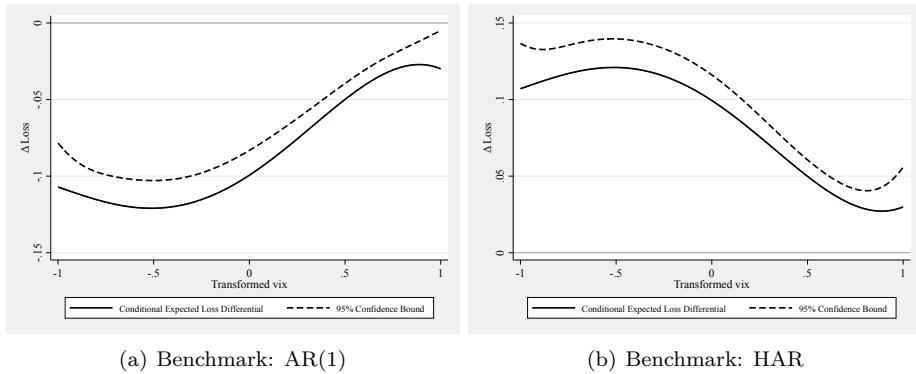


Figure 1. Diagnostic plots for one-versus-one CSPA tests

The results presented above are generated under the default setting of `cspa`, except that we set the Newey–West lag parameter at `lag(11)` to account for serial dependence following Li, Liao, and Quaedvlieg (2022). The significance level is  $\alpha = 5\%$ , which may be modified via the `siglevel()` option. By default, the series basis consists of  $m = 5$  series terms (that is, `m(5)`), corresponding with a fourth-order Legendre polynomial. In addition, the conditioning variable `vix` is transformed onto the  $[-1, 1]$  interval via the rank transformation (that is, `method(rank)`). Robustness checks with respect to these choices can be conveniently implemented by modifying the corresponding options, as illustrated in section 4.5 below.

### 4.3 One-versus-all CSPA test

We next demonstrate how to implement the CSPA test with multiple competitors. Recall that `rv.dta` contains six competing volatility forecasts. We may thus perform one-versus-all CSPA tests with one model as the benchmark and all the other five models as competitors, corresponding with  $J = 5$ . Under the null hypothesis, the benchmark model weakly dominates all the others uniformly across the conditioning space spanned by the VIX. The following command implements such a test, with HAR being the benchmark:

```
. cspa vix har ar1 ar22 ar22lasso harq arfima, lag(11) plot
```

Transformation on vix: Rank.

Benchmark	CSPA Test (5%)	p-value
har	-0.0095 (reject)	0.0002

From the table, we see that the CSPA null hypothesis for the HAR benchmark is strongly rejected. Although the one-versus-one test presented in section 4.2 suggests that HAR outperforms AR(1), the more stringent one-versus-all test reveals that HAR is no longer uniformly superior once additional competitors join the competition.

For comparison, we further implement the test for the HARQ benchmark as follows:

```
. cspa vix harq ar1 ar22 ar22lasso har arfima, lag(11) plot
```

Transformation on vix: Rank.

Benchmark	CSPA Test (5%)	p-value
harq	0.0054 (non-reject)	0.6778

Here we see that the CSPA null hypothesis cannot be rejected at any conventional significance level, offering formal statistical evidence for HARQ's superiority relative to the other competitors across conditioning states.

In the present setting with multiple competitors, `plot` draws the estimated lower envelope function  $\min_{1 \leq j \leq J} \hat{h}_j(\cdot)$ , which depicts the worst-case relative performance of the benchmark (in comparison with the “toughest” competitor) across the different conditioning states. The associated upper confidence bound  $\min_{1 \leq j \leq J} \{\hat{h}_j(\cdot) + n^{-1/2} \hat{k}_{1-\alpha} \hat{\sigma}_j(\cdot)\}$  is also plotted; if this bound ever falls below zero, the null hypothesis is rejected.

The left panel of figure 2 shows the plots associated with the HAR benchmark; the right panel, the plots associated with the HARQ benchmark. From the left panel, we see that the lower envelope function is always below zero, suggesting that at every conditioning state the HAR benchmark is outperformed by some competitor. The upper confidence bound dips below zero over the low-VIX region, which explains the CSPA test's rejection decision. On the other hand, we see from the right panel that the lower envelope function for the HARQ benchmark is generally “more positive”, indicating its better relative performance. The associated upper confidence bound is always above zero, which is why the null hypothesis is not rejected for HARQ.

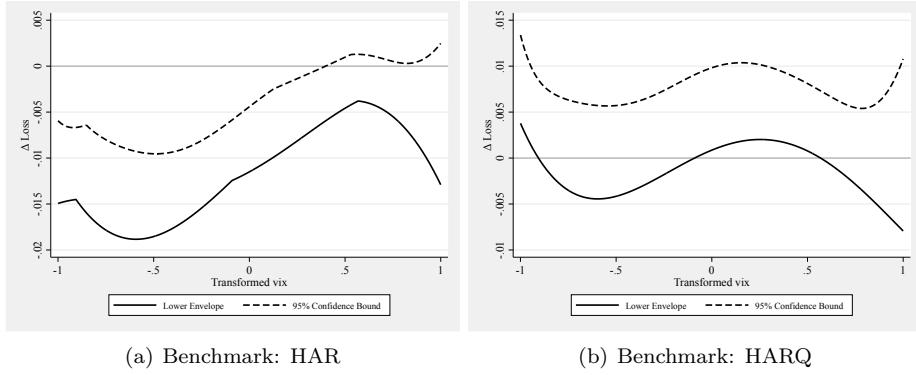


Figure 2. Diagnostic plots for one-versus-all CSPA tests

Note that for the one-versus-all test the `plot` option draws only the lower envelope function and its upper confidence bound, which concisely summarize the information from all  $J$  functional estimates. That noted, it may be useful for diagnostic purposes to also examine the individual  $\hat{h}_j(\cdot)$  estimates of conditional expected loss differentials. Their plots can be added by further calling the `detail(color)` option, as illustrated in figure 3.

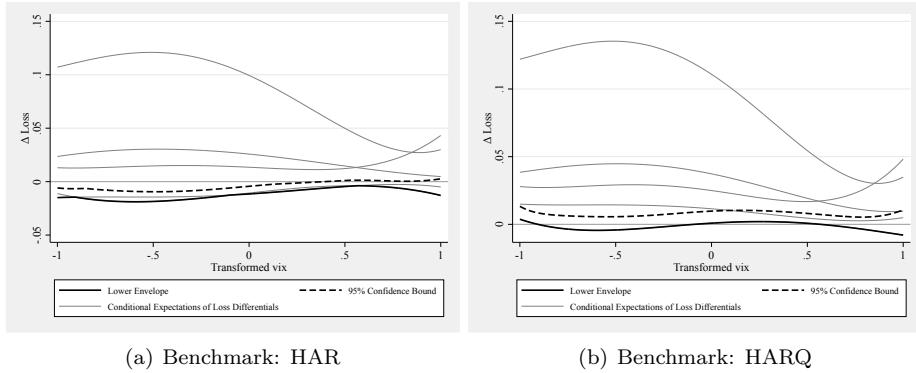


Figure 3. Detailed diagnostic plots for one-versus-all CSPA tests

Here the `detail(gray)` option draws the estimated conditional expectation functions of loss differentials,  $\hat{h}_j(\cdot)$ , for  $1 \leq j \leq 5$ , and the corresponding lower envelope function (the black solid curve). Because the lower envelope function consists of the minimum values of  $\hat{h}_j(\cdot)$ , that is,  $\min_{1 \leq j \leq J} \hat{h}_j(\cdot)$ , it always coincides with some parts of the  $\hat{h}_j(\cdot)$  curves. As a choice of design, we do not differentiate these gray curves using different line patterns so as to avoid a potential excess of information, especially for applications with many competing models. That being said, users may customize such

plots according to their specific needs using the more detailed output generated by the `excel` option.

The left panel of figure 3 reveals that, although HAR does not pass the CSPA test, its estimated conditional expected loss is lower than three competitors, as evidenced by the three  $\hat{h}_j(\cdot)$  curves above zero; but it is “beaten” by the other two competitors across all conditioning states. Note that the lower envelope function here is solely formed by these two superior competitors. Also, note from the right panel that the lower envelope for the HARQ benchmark coincides with one particular  $\hat{h}_j(\cdot)$  curve (overlaid by the lower envelope), which is associated with the ARFIMA model.

## 4.4 CSMS

As discussed in section 2.2, the CSMS may be constructed by collecting all benchmarks that are not rejected by the CSPA test. In particular, the one-versus-all tests performed in section 4.3 suggest that, at the 95% confidence level, HARQ belongs to the CSMS but HAR does not. We may further implement the test for the other benchmarks to decide whether they should be included in the CSMS.

The `cspa` command offers a more compact way to accomplish this task. With the `csms` option turned on, the command will automatically perform the requisite tests by rotating the benchmark role across all competing models, as shown below.

```
. cspa vix ar1 ar22 ar22lasso har harq arfima, lag(11) csms
```

Transformation on vix: Rank.

Benchmark	CSPA Test (5%)	p-value
ar1	-0.1156 (reject)	0.0002
ar22	-0.0174 (reject)	0.0002
ar22lasso	-0.0339 (reject)	0.0002
har	-0.0095 (reject)	0.0002
harq	0.0052 (non-reject)	0.6834
arfima	0.0057 (non-reject)	0.7876

The 95% CSMS = {harq, arfima}.

Here the output table reports the CSPA test statistics and the associated *p*-values for all six one-versus-all tests with different benchmarks. It also reports the CSMS, which in the current example consists of the HARQ and ARFIMA models. Note that when `csms` is active it is unnecessary to distinguish *benchmark* from *competitors* because all models under consideration are treated symmetrically.

## 4.5 Options

We highlight a few options that are more likely to be useful in typical empirical applications. For brevity, the illustration below focuses on the one-versus-all CSPA test for the HAR benchmark. Recall from section 4.3 that the null hypothesis is rejected when the testing procedure is implemented under `cspa`'s default setting. Our goal here is to perform some robustness checks with respect to this choice.

We first consider the number of series terms,  $m$ . Instead of using the default value, we may manually set it to 4, 6, 8, or 10 by modifying the `m()` option as shown below. The null hypothesis is still strongly rejected in all of these settings.

```
. cspa vix har ar1 ar22 ar22lasso harq arfima, lag(11) m(4)
```

Transformation on vix: Rank.

Benchmark	CSPA Test (5%)	p-value
har	-0.0102 (reject)	0.0002

```
. cspa vix har ar1 ar22 ar22lasso harq arfima, lag(11) m(6)
```

Transformation on vix: Rank.

Benchmark	CSPA Test (5%)	p-value
har	-0.0093 (reject)	0.0002

```
. cspa vix har ar1 ar22 ar22lasso harq arfima, lag(11) m(8)
```

Transformation on vix: Rank.

Benchmark	CSPA Test (5%)	p-value
har	-0.0087 (reject)	0.0002

```
. cspa vix har ar1 ar22 ar22lasso harq arfima, lag(11) m(10)
```

Transformation on vix: Rank.

Benchmark	CSPA Test (5%)	p-value
har	-0.0086 (reject)	0.0002

Next we recall that, under the default setting, the approximating functions are formed as Legendre polynomials of the rank-transformed `vix`. The purpose of using the rank transformation is to make the conditioning variable approximately uniformly distributed on  $[-1, 1]$ , which helps reduce the multicollinearity among the series regressors. We may also use the other types of transformations by modifying the `method()` option. Because the distribution of VIX is close to being log-normal, `method(lognormal)` is a natural choice for this task. From the implementation below, we see that this modification does not affect the rejection decision of the CSPA test.

```
. cspa vix har ar1 ar22 ar22lasso harq arfima, lag(11) method(lognormal)
```

Transformation on vix: Lognormal.

Benchmark	CSPA Test (5%)	p-value
har	-0.0085 (reject)	0.0002

Finally, we note that the diagnostic figures generated by the `plot` option are plotted on the transformed scale over the  $[-1, 1]$  interval. Alternatively, we may also set the horizontal axis to the original untransformed scale by using the `plotu` option. To clarify, the only effect of switching from `plot` to `plotu` is to compress or stretch the plotted curves along the horizontal axis, resulting in an alternative graphical presentation. The underlying testing results are completely unchanged. A concrete demonstration is given below, with the new plot displayed on figure 4. From the figure, we see that the statistical evidence against the null hypothesis mainly stems from the subregion with the VIX below 25%. We also note that the gap between the confidence bound and nonparametric estimate is fairly wide when the VIX is, say, above 50%. This is because the effective sample size for the local estimation on this region is relatively small, resulting in imprecise nonparametric estimates. This is also why the corresponding segment is “compressed” under the rank-transformed scale as previously shown on the left panel of figure 2.

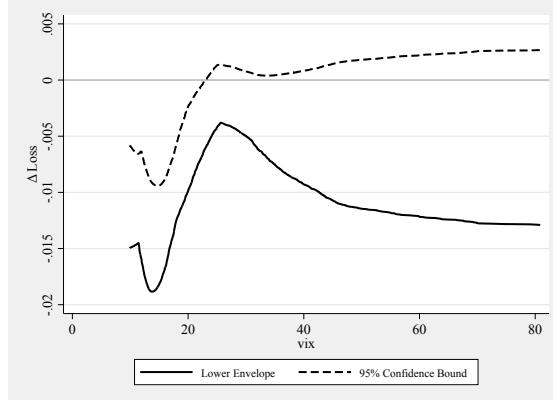


Figure 4. Diagnostic plot for the one-versus-all CSPA test generated by `plotu` (benchmark: HAR)

## 5 Conclusion

In this article, we described a new command, `cspa`, that conducts the CSPA test recently developed in Li, Liao, and Quaedvlieg (2022). Specifically, this command enables researchers to compare the performance of predictive methods across the conditioning space and construct the corresponding CSMS. The main functionalities of the `cspa` command were illustrated in an empirical application on volatility forecasting.

The proposed command is particularly useful when the researcher is interested not only in the competing models' average performance but also in their more detailed behavior across different conditioning states (for example, business cycle, financial conditions, and demographic characteristics). The method is general enough to accommodate time-series data with unknown forms of serial correlation and heteroskedasticity, and it can be readily specialized for applications with independent random samples under the default setting with `lag(0)`. An unconditional superior predictive ability test analogous to that proposed by Hansen (2005) may be implemented by setting `m(1)`, which is equivalent to conditioning on only the constant term. In principle, one can also use `cspa` to implement a nonparametric version of the CEPA test of Giacomini and White (2006) by treating a conditional moment equality as a pair of conditional moment inequalities. That being said, this task may be more directly accomplished in Stata via a nonparametric regression by using, for example, the `tssreg` command proposed in Li, Liao, and Gao (2020).

The current version of `cspa` may be improved in several ways. One is to allow for a multidimensional conditioning variable  $X_t$ . Because of the curse of dimensionality, how to "optimally" choose the approximating functions become a more relevant and challenging problem. The multidimensional setting would also require a more advanced visualization design than the simple plots of loss differential and confidence bound curves adopted in the one-dimensional setting here. It is also conceivable that other types of approximation basis functions may be more suitable than the Legendre polynomials in certain empirical settings, which may be a useful direction for future development.

## 6 Programs and supplemental materials

To install a snapshot of the corresponding software files as they existed at the time of publication of this article, type

```
. net sj 22-4
. net install st0696      (to install program files, if available)
. net get st0696         (to install ancillary files, if available)
```

## 7 References

Bollerslev, T., A. J. Patton, and R. Quaedvlieg. 2016. Exploiting the errors: A simple approach for improved volatility forecasting. *Journal of Econometrics* 192: 1–18. <https://doi.org/10.1016/j.jeconom.2015.10.007>.

Corsi, F. 2009. A simple approximate long-memory model of realized volatility. *Journal of Financial Economics* 7: 174–196. <https://doi.org/10.1093/jjfinec/nbp001>.

Diebold, F. X., and R. S. Mariano. 1995. Comparing predictive accuracy. *Journal of Business and Economic Statistics* 13: 253–263. <https://doi.org/10.1080/07350015.1995.10524599>.

Giacomini, R., and H. White. 2006. Tests of conditional predictive ability. *Econometrica* 74: 1545–1578. <https://doi.org/10.1111/j.1468-0262.2006.00718.x>.

Hansen, P. R. 2005. A test for superior predictive ability. *Journal of Business and Economic Statistics* 23: 365–380. <https://doi.org/10.1198/073500105000000063>.

Jann, B. 2005. moremata: Stata module (Mata) to provide various functions. Statistical Software Components S455001, Department of Economics, Boston College. <https://ideas.repec.org/c/boc/bocode/s455001.html>.

Li, J., Z. Liao, and M. Gao. 2020. Uniform nonparametric inference for time series using Stata. *Stata Journal* 20: 706–720. <https://doi.org/10.1177/1536867X20953576>.

Li, J., Z. Liao, and R. Quaedvlieg. 2022. Conditional superior predictive ability. *Review of Economic Studies* 89: 843–875. <https://doi.org/10.1093/restud/rdab039>.

White, H. 2000. A reality check for data snooping. *Econometrica* 68: 1097–1126. <https://doi.org/10.1111/1468-0262.00152>.

### **About the authors**

Jia Li is the Lee Kong Chian Professor of Economics at the School of Economics at Singapore Management University.

Zhipeng Liao is an associate professor of economics at the Department of Economics at the University of California–Los Angeles.

Rogier Quaedvlieg is an economist at the Monetary Analysis Division in the Directorate Monetary Policy of the European Central Bank in Frankfurt, Germany.

Wenyu Zhou is an assistant professor at the International Business School at Zhejiang University.