

Discussion of “Feature Matching in Time Series Modeling” by Y. Xia and H. Tong

Bruce E. Hansen

1. INTRODUCTION

Xia and Tong have written a provocative and stimulating paper. Among the many topics raised in their paper, I would like in particular to endorse several of their postulates:

1. All models are wrong.
2. Observations are not error-free.
3. Estimation needs to account for the above two issues.

As described in the paper, suppose that we observe a process $\{y_t : t = 1, \dots\}$ for which we have a model $\{x_t(\theta) : t = 1, \dots\}$ which depends upon an unknown parameter θ . Let $F_x(\theta)$ denote the joint distribution of the $x_t(\theta)$ process and F_y the joint distribution of the observables. When we say that the model is wrong, we mean that there is no θ such that $F_x(\theta) = F_y$. If we think of the distribution F_y as a member of a large space of potential joint distributions, then the set of joint distributions $F_x(\theta)$ constitutes a low-dimensional subspace of this larger space. While there is no true θ , we can define the pseudo-true θ as the value which makes $F_x(\theta)$ as close as possible to F_y . This requires specifying a distance metric between the joint distributions

$$d(\theta) = d(F_x(\theta), F_y)$$

and then we can define the best-fitting model $F_x(\theta)$ by selecting θ to minimize $d(\theta)$. The relevant question is then: what is the appropriate distance metric?

2. CATCH-ALL ESTIMATION

Xia and Tong recommend what they call a “catch-all” approach, where the distance metric is a weighted sum of squared k -step forecast residuals. They show that in some situations this criterion allows consistent estimation of the parameters of the true latent process.

Bruce E. Hansen is Professor, Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, Wisconsin 53706, USA (e-mail: behansen@wisc.edu).

Their Theorem C requires that the latent process is deterministic, but the result might hold more broadly.

This can be illustrated in a very simple example of a latent AR(1) with additive measurement error. Suppose that the latent process is

$$x_t = \theta x_{t-1} + \varepsilon_t$$

and the observed process is

$$y_t = x_t + \eta_t,$$

where ε_t and η_t are independent white noise. In this case, it is well known that y_t has an ARMA(1, 1) representation

$$(1) \quad y_t = \theta y_{t-1} + u_t - \alpha u_{t-1},$$

where u_t is white noise and $0 \leq \alpha < 1$.

Xia and Tong propose estimation based on k -step forecast errors. The k -step forecast equation for the observables is

$$(2) \quad y_{t-1+k} = \theta^k y_{t-1} + e_t(k),$$

where

$$e_t(k) = \sum_{j=0}^{k-1} \theta^j (u_{t+k-j-1} - \alpha u_{t+k-j-2}).$$

Xia and Tong’s estimator is based on a weighted average of squared forecast errors. For simplicity, suppose all the weight is on the k th forecast error. The estimator is

$$\hat{\theta}_{\{k\}} = \arg \min_{\theta} \sum_{t=1}^T (y_{t-1+k} - \theta^k y_{t-1})^2$$

which has the explicit solution

$$\hat{\theta}_{\{k\}} = \left(\frac{\sum_{t=1}^T y_{t-1} y_{t-1+k}}{\sum_{t=1}^T y_{t-1}^2} \right)^{1/k}.$$

We calculate that as $n \rightarrow \infty$

$$\hat{\theta}_{\{k\}} \xrightarrow{p} \theta_{\{k\}} = \theta(1 - c)^{1/k},$$

where $c = \alpha \sigma_u^2 / \theta \sigma_y^2$, $\sigma_u^2 = E u_t^2$ and $\sigma_y^2 = E y_t^2$.

Thus for any k , $\hat{\theta}$ is inconsistent as an estimator of θ . But as k gets large the discrepancy gets smaller, as $(1 - c)^{1/k} \rightarrow 1$ since $c < 1$. Thus as $k \rightarrow \infty$

$$(3) \quad \theta_{\{k\}} \rightarrow \theta.$$

This derivation assumed that the estimator is based on the k th forecast error, but it extends to the case of a weighted average.

The convergence (3) is an extension of Xia and Tong's Theorem C. It shows that estimation by minimizing the squared k -step forecast residual is consistent for the parameter of the latent AR(1), as k is made large.

One trouble with this approach is that the estimator is quite inefficient. We can calculate that

$$T \text{var}(\hat{\theta}_{\{k\}}) \simeq \left(\frac{1}{k\theta^k} \right)^2 \rightarrow \infty$$

as $k \rightarrow \infty$. This means that the variance of the Xia–Tong estimator is increasing in k (and unbounded). This is especially troubling since the parameters of (1) can be estimated by standard ARMA methods. The implication is that while the catch-all approach has some useful robustness properties, there is no reason to expect the estimator to be efficient.

3. MEASUREMENT ERROR AND NONPARAMETRIC IDENTIFICATION

Xia and Tong emphasize that measurement error is empirically relevant and time series methods should take it seriously. While I agree, we also need to acknowledge that measurement error raises many troubling problems. Of primary importance, I believe, is the vexing issue of nonparametric identification—whether the parameters of interest are uniquely determined by the distribution of the observables. As is known from the random sampling context, measurement error complicates identification. In general, additional information or structure is required to identify the parameters of an unobserved latent process. It is not sufficient to simply introduce a new estimator.

We can see this quite simply by examining the spectral density. Suppose as above that x_t is the process of

interest and the observed process is $y_t = x_t + \eta_t$ where η_t is i.i.d. measurement error with variance σ_η^2 . Letting $f_x(\lambda)$ and $f_y(\lambda)$ be the spectral densities of x_t and y_t , we know that

$$f_y(\lambda) = f_x(\lambda) + \sigma_\eta^2.$$

The distribution of the observables y_t identifies $f_y(\lambda)$, but $f_x(\lambda)$ is not identified from knowledge of $f_y(\lambda)$ alone. Under the realistic assumption that σ_η^2 is unknown, $f_x(\lambda)$ can only be identified by knowledge of the structure of x_t [e.g., by knowing that x_t is an AR(1) as in the example of the previous section]. But if we acknowledge that our models for x_t are misspecified, we should view the true $f_x(\lambda)$ as nonparametric and hence without structure. It follows that the spectral density $f_x(\lambda)$ is not nonparametrically identified, and thus neither is the autocorrelation structure of x_t .

Nevertheless, some features are identified. While the spectral density is not point identified, it is interval identified. Let $\bar{f} = \min_\lambda f_y(\lambda)$. Observe that

$$f_y(\lambda) - \bar{f} \leq f_x(\lambda) \leq f_y(\lambda).$$

The two bounds are identified from $f_y(\lambda)$, so the spectral density $f_x(\lambda)$ of x_t can be bounded within this interval. The width of the interval is $\bar{f} = \min_\lambda f_x(\lambda) + \sigma_\eta^2$, which is thus an upper bound for the measurement error variance σ_η^2 .

What is particularly interesting is that while the level of $f_x(\lambda)$ is not identified, many of its most important features are identified, specifically, the peaks and troughs. What this means is that while full knowledge of the x_t process is not possible, important features can be identified from the distribution of the observables y_t . Knowledge of which features are identified in the presence of measurement error and/or misspecification helps focus attention on what can be learned about unobserved processes from observational data.

ACKNOWLEDGMENTS

Research supported by the National Science Foundation.