SIMS MACRO GENERAL EXAM QUESTION

Suppose we have data $\{y_1, \ldots, y_T\}$ from a stochastic process that we know is the sum of two independent first order autoregressive underlying processes. That is

$$\begin{aligned} x_t &= \rho x_{t-1} + \varepsilon_t \\ z_t &= \alpha z_{t-1} + \nu_t \\ y_t &= x_t + z_t \,. \end{aligned}$$

Assume $|\alpha| \leq 1$ and $|\rho| \leq 1$, ε and ν i.i.d. normal and independent of each other, and with unknown variances σ^2_ε and σ^2_ν.

(a) From data on y alone, can we estimate the four unknown parameters ρ, α, σ^2_ε and σ^2_ν consistently? Are there particular parameter values that create identification problems?

The autocovariance function of y will be the sum of the two ACF's for x and y, therefore of the form (using the usual abuse of notation in which the ACF is represented as the coefficients of a polynomial in the lag operator)

$$\frac{\sigma^2_\nu(1+\rho^2) + \sigma^2_\varepsilon(1+\alpha^2) - (\sigma^2_\nu \rho + \sigma^2_\varepsilon \alpha)(L + L^{-1})}{(1+\rho^2 - \rho(L+L^{-1}))(1+\alpha^2 - \alpha(L+L^{-1}))} \,.$$

Because the numerator is second order and symmetric in L and L^{-1}, while the denominator is 4th order, and also symmetric in L and L^{-1}, this is the ACF of an ARMA(2,1), i.e. second-order in its autoregressive component and first order in its MA component. So long as there is no cancellation between numerator and denominator, the coefficients of an ARMA can be estimated consistently. Since the ACF denominator has two roots on or outside the unit circle, they are the the roots of the AR operator in the MA process, and they can be estimated consistently (even if one or both are one). The numerator of the ACF has two coefficients, the constant and the coefficient on $(L + L^{-1})$. With ρ and α determined by the denominator, this gives enough equations, at least, to determine σ^2_ε and σ^2_ν. A really thorough answer might have checked that the linear mapping between the numerator ACF coefficients and the σ^2 parameters was non-singular. The mapping is

$$\begin{bmatrix} 1+\rho^2 & 1+\alpha^2 \\ -\rho & -\alpha \end{bmatrix} \begin{bmatrix} \sigma^2_\nu \\ \sigma^2_\varepsilon \end{bmatrix} = \begin{bmatrix} c_0 \\ c1 \end{bmatrix} \,,$$

where c_0 and c_1 are the two numerator ACF coefficients. It is easy to see that the left-hand side coefficient matrix is singular if $\rho = \alpha$, and not too hard to show (solving a quadratic equation) that it is singular only in that case. So the whole model is identified unless $\rho = \alpha$. When $\rho = \alpha$ the y process is just an AR(1) with root ρ, so the relative contributions of the σ^2_ε and σ^2_ν components of residual variance can't be identified.

We know that ARMA models generally can have identification problems from root cancellation. Could that happen here? That would require that the numerator of the ACF of y have a root that matches ρ or α. But the numerator polynomial is the weighted sum of two polynomials, one that has ρ as a root and one that has α as a root. That means the whole numerator can't have either ρ or α as a root, unless, again, $\rho = \alpha$.

(b) Can every ARMA(2,1) process be represented as such a sum of independent AR1's? Explain your answer.

Probably the easiest counterexample comes from the fact that the denominator of an ARMA(2,1) could have a pair of complex roots, while no real-valued univariate AR(1) can have a complex root.

(c) Describe a procedure for Bayesian sampling from the posterior distribution of the four parameters given data on y alone. Consider first the case where it is known that $|\rho|$ and $|\alpha|$ are less than one, then the case where it is possible that one or both of these parameters is one. Include discussion of what might be a reasonable prior and whether inference is likely to be sensitive to the prior.

In the stationary case, one approach would be to use the ACF to populate a covariance matrix Σ for y and then form the standard Gaussian log likelihood

$$-\tfrac{1}{2}\,|\Sigma| - \tfrac{1}{2}\vec{y}'\Sigma^{-1}\vec{y}\,.$$

This should form a well-behaved posterior density, and one could use, if no problem-specific prior information was available, flat priors on (-1,1) for α and ρ and, for example, flat priors on the logs of the σ^2 parameters. One would maximize the likelihood to get an initial estimate of the parameters, use the inverse second derivative matrix of the log likelihood, multiplied by about .25, as minus the covariance matrix of a jump distribution, and apply random-walk Metropolis posterior simulation.

One could also treat this as a state-space model with $y_t = x_t + z_t$ as the observation equation and the two autoregressive equations as the plant equations. In the stationary case again it would be reasonable to treat the initial values of x and z as drawn from their stationary joint distribution, with the same priors discussed above. The Kalman filter could then be used to evaluate likelihood for any given set of parameter values. MCMC calculations would work as above.

When one or both roots might be one, the main difference is that the initial conditions can no longer be given the (nonexistent) stationary marginal distribution. Conditioning on observed initial conditions, which is possible but not usually desirable in ordinary AR models, is not much help here, since the z and x initial states are still uncertain given the observed initial y. Subject matter considerations would be important here, but one approach that could work is to give the intitial x_1 z_1 values a $N(0, \mathrm{diag}(\sigma_x^2, \sigma_z^2))$ distribution, with σ_x^2 and σ_z^2 both large and fixed (i.e. not estimated). However this would create the common problem of allowing explanation of low frequency patterns in the data as predictable from initial conditions. So it would need to be combined with some prior guarding against this, e.g. a "Minnesota" dummy observations prior.