Comments on Model Comparison

Christopher A. Sims
Princeton University
sims@princeton.edu

September 10, 2010

Summary

- An interesting modeling approach, combining a very parsimonious regime-switching dynamies, a richly parameterized family of conditional distributions for returns, and an assumption of constant expected returns.
- The model is estimated by an ad hoc procedure that is not motivated and makes measures of fit difficult to interpret.
- There are no comparisons of model performance to what would seem to be its real competitors — GARCH and stochastic volatility models.
- The model structure could be enriched without explosion of the parameter space, allowing interesting tests of fit.

The estimation criterion

- Why does it make sense to penalize failure to match the ACF beyond the implicit penalty in the likelihood?
- What is the cost in likelihood? In what aspect of fit do we see the cost?
- Are the "likelihoods" in the tables with or without the penalty term? (Must be without, since more parameters makes fit worse?)

How persistent are the estimated states?

- The estimates show very persistent states. The sum of the mean durations in the two states for the authors' favored estimates is over a year.
- Figure 6 on p.13 shows a high probability of over 40 transitions between states during a two-year span, whereas the model's expected number of transitions would be less than four.

How persistent are the actual states?

- Figure 7 shows 200-day moving averages of volatility. The high volatility episodes show up as humps of width pretty close to 200 days. This suggests quite *non*-persistent episodes of high volatility.
- It appears that the ACF penalty is forcing the model to long persistence, when likelihood would suggest something closer to a simple mixture model.

Distribution shapes

- One of the states is labeled “good”, the other “bad”, on the basis of which times they seem to prevail. But we would like to see the shapes of the distributions to understand what is “good” and “bad” about them. The way the parameters of these distributions translate into distribution shapes is not familiar to all readers (e.g. me).
- It appears that the estimated distributions are just slightly leptokurtic and that the biggest difference between the two states is simply in the spread of the two distributions.

Generalized Hyperbolic "bad" pdf, with normal

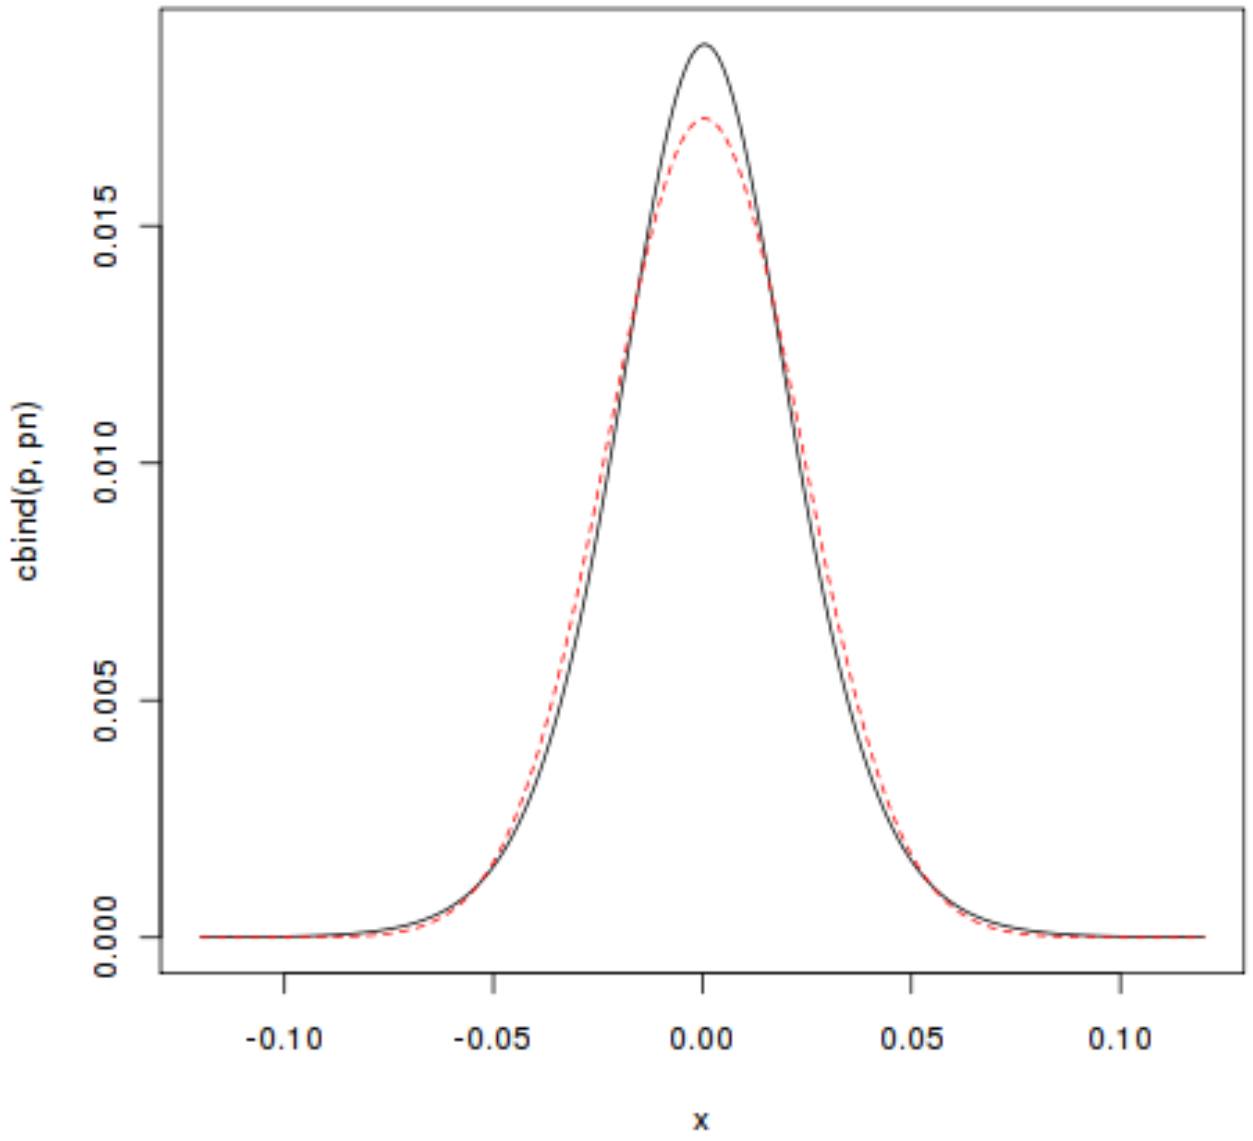

Generalized Hyperbolic "good" pdf, with normal

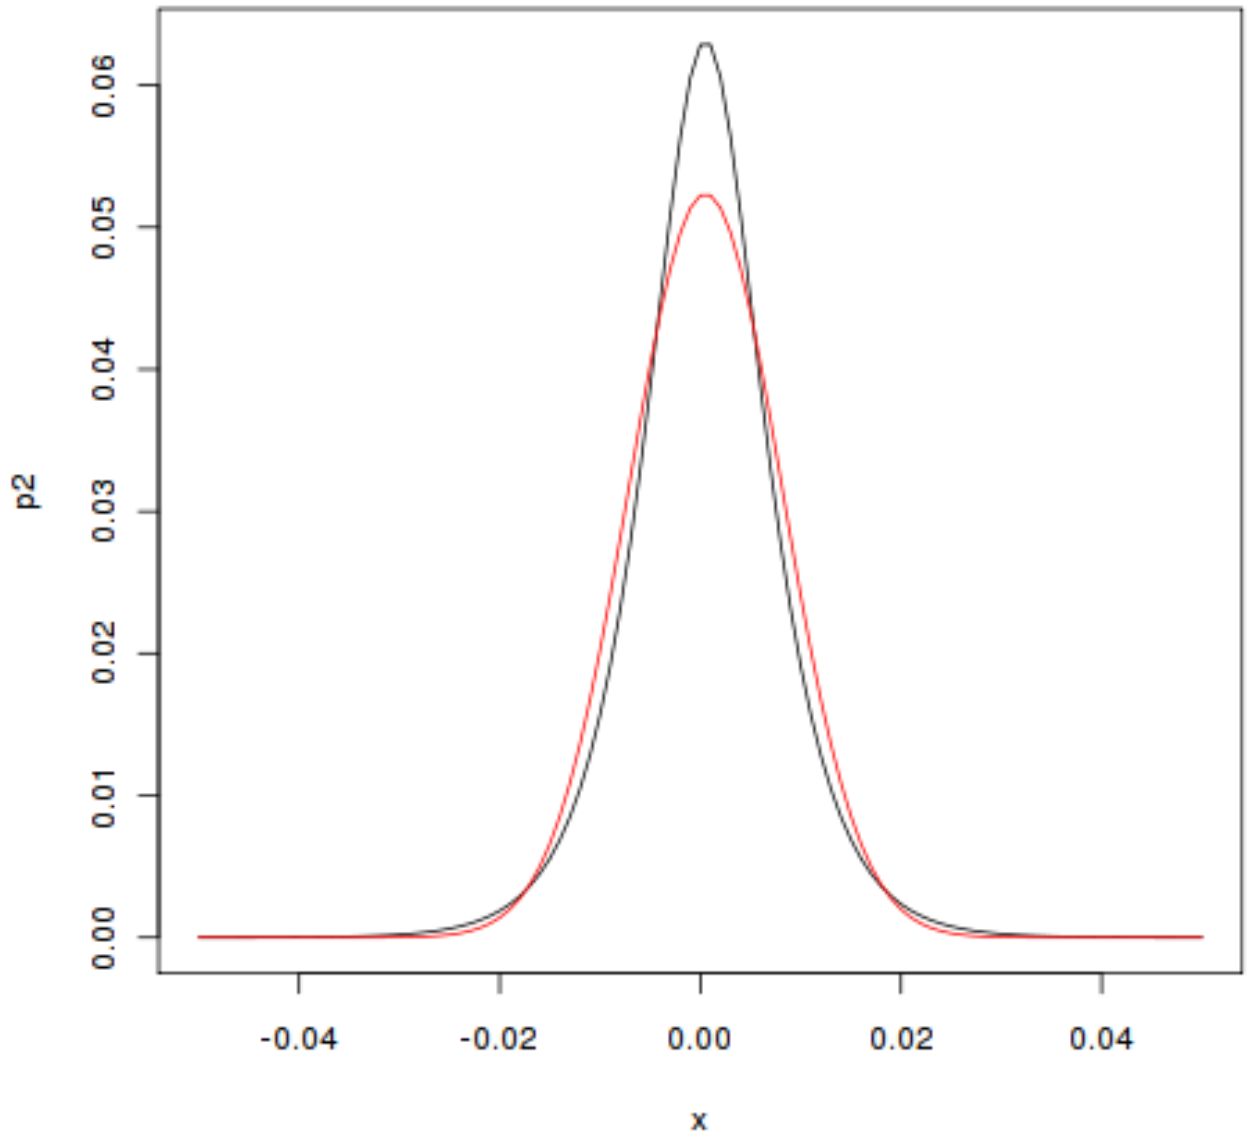

Things to do

- Why not k ordered states, with transitions concentrated along diagonals? This could approximate a smoother stochastic volatility model.
- Try relaxing the common-mean assumption. We don't expect to see much predictability, but we would be interested in how much room the data leaves for predictability. Not a test of a null hypothesis, a measure of uncertainty.
- Use priors, get posterior odds vs. stochastic volatility and ARCH. Matching unconditional distribution with K-S stat is weak test of fit.