Discussion of J.H. Stock and M. W. Watson "Phillips Curve Inflation Forecasts"

Adrian Pagan
University of New South Wales
and Queensland University of Technology

June 4, 2008

Stock and Watson provide a thorough (one might say exhaustive and probably exhausting) review of the forecasting performance of many inflation models. These include models with some economics in them and others that are purely statistical. Overall, the best of the statistical models seems to be their UC-SV model. This is a two component model of the form

$$\begin{aligned} \pi_t &= z_t^* + v_t \\ z_t^* &= z_{t-1}^* + u_t. \end{aligned}$$

As written z_t^* represents a permanent component to inflation and v_t is a transitory one. The shocks to the components are taken as having time varying stochastic volatilities (SV) σ_{jt}^2 of the form ln $\sigma_{jt}^2 = \ln \sigma_{jt-1}^2 + \eta_{jt}$. Consequently these have a unit root. The variances of the shock volatilities η_{jt} are equal. This would mean that the true volatilities would also be equal if the initial conditions were the same. The estimated ones can however vary as they depend upon the data.

Statistical models are generally judged by how well they fit and forecast over limited sample sizes and horizons. That is probably just as well here since it would be hard to think about targeting an inflation rate that really behaved like the UC-SV model, since it has inflation following a unit root and the variances of v_t and u_t being unbounded i.e. there are no second moments for the change in inflation. If you try to simulate a process like UC-SV it blows up very quickly.

I guess I am rather doubtful about whether I want to use a model like this, even if it produces good forecasts, as it would be hard to believe that we could keep inflation in a target range for very long if it behaved in such a way. It is only over short periods that this model makes much sense. Given this, I wonder why SV set up the UC model with a permanent shock. One could have chosen a very persistent process such as $z_t^* = .99z_{t-1}^* + u_t$ i.e. why choose 1.0 as the autoregressive coefficient rather than .99. Since putting any number in is arbitrary one might as well use something like .99 as that ties in better with what we hope the nature of the inflation process is. We are not estimating any extra coefficients and it seems highly likely that over short forecast horizons the forecasts using both sets of parameters would be very close. It might be different if we looked at an 8 period horizon since most central banks forecast both one and two years ahead.

SW conclude that there are periods of time when the UC-SV model can be beaten with models featuring economic variables, principally when there are large departures from the NAIRU or when one is in an extreme recession. But for most of the time the economic variables don't contribute much to forecasts. Of course there are good reasons why we like to think in terms of the influence of economic variables even if we cannot detect a precise role for them in forecasts. It is a well known fact that relatively statistical simple models win forecasting competitions. But, in a world in which we are increasingly forced to explain policy actions, any forecasts underlying them, and the risks associated with the latter, statistical model forecasts clearly have weaknesses. In practice many central banks use a Phillips type equation to give a central forecast - and even judgemental forecasts often have this as a base- and relegate the statistical model forecasts to the role of checking and "tweaking" the central forecast. It is hard to imagine any central banker not putting some faith in the role of excess demand in accounting for inflation outcomes even if measures of excess demand do little for forecasting inflation a year ahead. There are many reasons for why we often see this. The excess demand may need to be sustained for a long time, many measures of it are exceedingly volatile, inflation itself can have a lot of noise, it may only be when we exceed a threshold that there are serious inflation effects etc. These are all hard to capture with the data sets we typically face. But at some point economic variables need to enter the forecast otherwise it is not going to be easy to explain any actions you take as a consequence of the forecasts.

So let us remind ourselves why the UC-SV model might win the SW forecast competition. To do this assume that there is no SV in the UC model and that the variances of the two shocks u_t and v_t are in a fixed ratio. Then, it is well known that the $E_t\pi_{t+1}$ from the UC model (provided $E_t(v_{t+1}) = 0$) is[1]

$$E_t\pi_{t+1} = (1-q)\sum_{j=0}^{\infty} q^j \pi_{t-j},$$

where the parameter q is determined solely by the ratio of the variances of v_t and u_t. This is the exponentially weighted moving average (EWMA) forecasting formula that is widely used in industry for forecasting the level of product demand. So perhaps it is not surprising that it produces a forecast that is good for inflation. It has also been used recently in the financial literature to forecast series that are random walks in which the drift term changes over time i.e. it has some robustness to breaks in those series. So it may be a good model for inflation as well.

To develop this theme further we note that the parameter $q = -\alpha$ where α is the MA coefficient in the ARIMA representation of the UC model i.e. $\Delta\pi_t = \varepsilon_t + \alpha\varepsilon_{t-1}$. SV note that α has been varying a good deal over US history and, recently, $\alpha \cong -.85$, and so one would expect a small weight to be put on past inflation. Indeed a large negative value of α is really consistent with inflation not having a unit root, and the way this shows up in the UC model is for the variance of the transitory shocks to become large relative to the permanent ones.[2] It should be noted for later reference that the MLE of the MA(1) coefficient values in recent times is hard to distinguish from other values (based on the likelihood). Thus the MLE of -.76 found by fitting an MA(1) to the change in inflation of the GDP deflator inflation over the period 2001/1-2008/1 is hard to distinguish from a range of other values - see figure 1

Now one reason the EWMA has been popular is that it represents a simple forecasting mechanism that is relatively robust to structural change, provided one modifies q at different times. A plot of the inflation rate suggests that such changes have occurred, and there is an extensive literature maintaining that such breaks have occurred in many countries in the past three decades,

[1] SW look at predicting inflation over the next year but we look at one step here for exposition.

[2] Of course the exponentially weighted MA forecast above comes from the Kalman Filter and that is not really appropriate when the conditional variances vary and are unobservable. Therefore the actual forecasts SW would produce would be non-linear functions of the data.

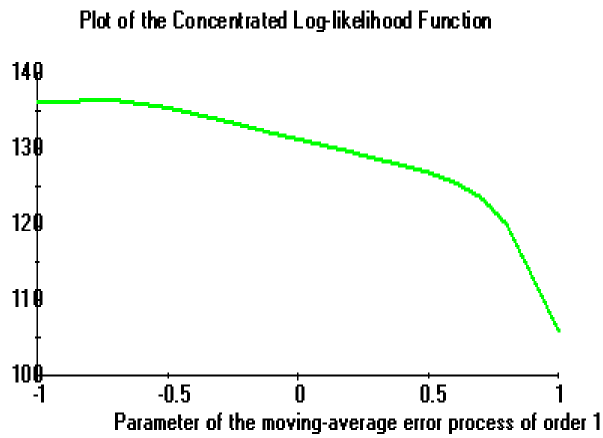

as well as in the U.S. The key question then becomes one of how to vary q in response to such developments and this is done with the SV part of SW's UC-SV model, since the estimated relative volatilities can change over time, leading to a change in q.

So this leads us to ask what the implications of the SW paper are for forecasting. This may not be a fair question since their brief seems to have been to ask if Phillips curve type economic variables are useful for forecasting rather than to ask what is the best forecasting method. So I am possibly being unfair when I ask if their UC-SV model is the best forecasting mechanism, but I think their results are sufficiently striking for me to make some comments on this.

As mentioned above, in SW's case q adapts to the data to account for breaks through the relative size of the stochastic volatilities. Are there other ways of doing this? Pesaran and Timmermann (2007) point out that we need to detect when a break in the inflation process took place and also the size of the break i.e. we need to know when to vary q and by how much. There has been much research on methods to do the former but less on the latter. However there is is now an emerging literature on it. Pesaran and Timmermann (2007) have proposed the idea that one should average forecasts not across models (SW show that this didn't give much advantage) but across the windows over which parameter estimation is performed prior to making the forecast. They demonstrate that this can yield improvements in forecasting a random walk process in the presence of breaks. Pesaran and Pick (2008) show that there are theoretical reasons to expect that this procedure will improve forecasts in the presence of breaks. They also look at EWMA forecasts for different q values and then average the forecasts from these. An advantage of focussing upon the EWMA approach is that no judgement is being made about the nature of the inflation process (one still uses a weighted average of inflation rates). If the inflation series was white noise than one would simply put $q = 0$. So I think it would be interesting to compare the forecasts from this methodology with those from the UC-SV model.

Another method of locating and sizing the breaks is the indicator saturation approach of Hendry et al. (2004). The indicators are dummy variables indicating where breaks took place and these are selected by an automatic model selection device. Hendry et al show that, despite the large number of possible dummies, this turns out to be a very efficient way of learning about multiple breaks. They point out that people are essentially doing this when they perform recursive estimation, since the extra observation is equivalent to using a dummy variable. Once the breaks are found the magnitude should be the values of the dummy variable coefficients.This has not yet been applied in a forecasting context but was used by Castle and Hendry(2008) to isolate breaks in the UK inflation process. Unlike the Pesaran-Timmermann method that assumes there is a unit root in the breaking process, the indicator variable saturation method makes no such restrictions, and so this is a potential advantage of it.

Now let us think about introducing economic information into the forecasts. Traditionally this involved replacing z_t^* with some functions of lagged inflation, unemployment, supply side effects, deviations of this from the NAIRU, expectations etc. In the US literature this is often referred to as Gordon's triangle model, although models like it have been present in many countries since the 1970s. Whatever is introduced needs to be combined together to produce a persistent component. The flaws in the approach are the need to estimate the parameters in any such relation as providing forecasts of these covariates. Consequently, it is possible that we will do worse than a model that ignores the effects even if we believed that the economics in them tells us something about what might have driven an observed inflation path. That will almost certainly happen if the parameters are imprecisely determined, and one would have to say that this is indeed true of most Phillips curves. Moreover many of the variables added into the relation can change quite dramatically as a result of data revisions. So even if the forecasts are good with data that has been finally revised, the need to use real-time data may result in the forecasts being quite poor - see for example Robinson et al (2003) for an Australian example. Since SW didn't use real time data it would seem that the Phillips curve based forecasts they report would most likely be better than they would be in real time. It is only with extreme movements in the determining variables - very high unemployment relative to the NAIRU or expectations relative to say the target- that we can observe big enough effects on inflation to offset these difficulties. However it should be noted that it may be possible to use economic information to reliably signal the direction of change in inflation, as found in Robinson et al (2003) for Australia and Fisher et al (2002) for the US,and in many contexts this might well be sufficient.

An alternative mechanism is to generalize the UC model to allow the economic variables to influence either Δz_t^* or the transitory component. Thus we might write

$$\begin{aligned} \pi_t &= z_t^* + x_t'\gamma + v_t \\ \Delta z_t^* &= x_t'\delta + u_t \end{aligned}$$

In this case the forecast of π_{t+1} is an EWMA of $\pi_t - x_t'\gamma$ augmented with $x_t'\delta$. A difficult with this form is that $x_t'\delta$ can't be very persistent if u_t is not to be highly serially correlated. One's feeling is that it would be hard to gain much for forecasting from such an extension.

There are other non-UC models that have been developed for various countries. One of these, done by us for Australia and appearing as Gruen et al (1999) in the JME issue on the Phillips Curve - was

$$\Delta\pi_t^4 = x_t'\delta + \beta(\Delta\pi_{t-1} - \Delta\pi_{t-4}) + e_t,$$

where $\Delta\pi_t^4$ is annual inflation, x_t had information on inflation expectations, the stochastically NAIRU gap and import price inflation. This is related to the Atkeson-Ohanian (AO) model that SW discuss and which is probably the main competitor to UC-SV. In AO $\beta = 0, \delta = 0$. Now the equation was written as above since we found that lagged quarterly inflation had a role in determining yearly inflation i.e. it was not enough to use just π_{t-1}^4 in predicting π_t^4. I briefly looked at this in the US data case and it seemed that terms like π_{t-1} and π_{t-4} entered the equation using US data.

Figure 2: q (proxied by negative of AR(1) coefficient estim over rolling 10 year periods) and HP output Gap

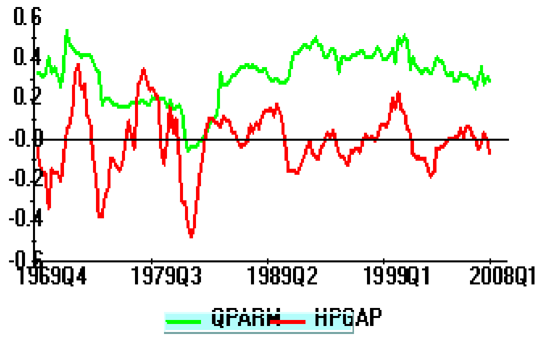

Finally we might just make q a function of some economic variables and so change the exponential weights. To get some idea of whether this would work we need a series on q_t. Ideally we would want to fit an MA(1) to the change in inflation, but due to time constraints I estimated q in a simple way by fitting an AR(1). We use a 10 year rolling horizon and so get a sequence of q values starting from the 1970s. There are problems with this. The AR(1) coefficient ρ is related to the MA coefficient by solving a quadratic equation. This yields a highly non-linear relation. So a value of the MA coefficient of -.76 becomes an AR(1) coefficient of -.47 while a value of around -.2 is much the same for the MA. Thus if the likelihood is flat in the MA(1) coefficient - as was seen in figure 1- then we can easily get a value of the AR(1) coefficient which is well away from what the MA(1) coefficient might be.

Figure 2 hows the series on q_t obtained with this method. The graph is best at identifying changes rather than the precise values e.g. the MA(1) coefficient has become larger in absolute terms in recent years and that has meant a larger value for q_t. This figure also shows how the output gap (found with the HP filter) has varied. Figure 3 shows a corss plot of q against the gap.

There are clearly times when q changes quite markedly. The HP filtered output gap also seems to have an influence upon q which tends to be in agreement with what SW find i.e. it is mainly when the output gap gets large in a recession that the economic variables are important. It is possible that one might find it useful to build models of q in this way in order to perform forecasts which introduce economic influences such as output gaps and unemployment.

References

Castle, J.L. and D.F. Hendry (2008), " The Long-Run Determinants of UK Wages 1860-2004", unpublished paper, Department of Economics, Oxford University

Gruen, D, A.R. Pagan and C. Thompson (1999), "The Phillips Curve in Australia", *Journal of Monetary Economics*, 44, 223-258.

Hendry, D.F., S. Johansen and C. Santos (2004), "Selecting a Regression Saturated by Indicators", unpublished paper, Department of Economics, Oxford University

Fisher, J.D.M., C.T. Liu and R. Zhou (2002), "When Can we Forecast Inflation", *Federal Reserve Bank of Chicago Economic Perspectives*, 1st Q, 30-42

Pesaran, M.H. and A. Timmermann (2007), "Selection of Estimation Window in the Presence of Breaks", *Journal of Econometrics*, 137(1), 134-161.

Pesaran, M.H. and A. Pick (2008), "Forecasting Random Walks Under Drift Instability", unpublished paper, Cambridge University

Robinson, T., A. Stone and M. van Zyl (2003), "The Real-Time Forecasting Performance of Phillips Curves", *Reserve Bank of Australia Discussion Paper 2003-12*