Non-rational expectations and the transmission mechanism

Summary of Working Paper no. 448 Richard Harrison and Tim Taylor

Models are important tools that economists use to help them understand the behaviour of the economy. Many macroeconomic models assume that the decisions of households and firms should depend on their expectation of future events. For example, a household's saving decision is likely to be influenced by an assessment of the income that is expected to be earned in the future. And the price a firm decides to set for its product is likely to depend on its view of the costs of production that it will incur over the period until it next resets its price. An important assumption for such models is how households and firms form their expectations of future earning and costs. The dominant assumption in macroeconomics is that expectations are formed in a way that is 'rational' (or 'model consistent'). An implication is that expectations are correct on average and that the difference between expected and actual outturns is unpredictable. In other words, households and firms do not make persistent mistakes when predicting future earnings or costs.

The rational expectations assumption is a very strong one, implying that households and firms have a lot of information about the structure of the economy. This has led economists and policymakers to examine the effects of alternative 'non-rational' expectations assumptions. Relative to the benchmark assumption of rational expectations, models that include non-rational expectations face two challenges. The first is the need to specify the mechanism through which expectations are generated. The second is how to capture the way that expectations of future earnings and costs affect the decisions that households and firms make about their current savings and pricing.

This paper is concerned with the second challenge. There are two main alternatives to modelling decision making when expectations are non-rational. To see the difference between these, suppose, as an example, that a household makes a decision over how much to save and how much to spend. The decision depends on the household's expectations of future earnings: higher future earnings allow the household to borrow to finance higher spending today. There are two ways to characterise how the household decides how much to spend and save.

The first approach relies on the consumption 'Euler equation', which states that the household's current consumption should depend on the expected level of consumption next period and the real interest rate. Other things equal, a higher real interest rate will encourage households to consume less and save more. This approach to non-rational expectations therefore assumes that household consumption is determined by the Euler equation, but with a non-rational expectation of future consumption. The second approach is to characterise the household's consumption decision in terms of the household's expectations of its entire lifetime income. Other things equal, the higher the household's expected lifetime income, the higher the household's current consumption. In this approach, consumption is therefore determined by non-rational expectations of lifetime income. We call this the 'long-horizon' approach.

Under rational expectations, the 'Euler equation' and 'long-horizon' approaches give identical answers: the household's consumption is the same in both cases. But under non-rational expectations, the predictions for consumption can be different. The purpose of this paper is to investigate how significant these differences may be. To do so, we build a model of household and firm behaviour under three assumptions: rational expectations, non-rational 'Euler equation' expectations and non-rational 'long-horizon' expectations. We then compare the behaviour of key variables for these variants of the model.

We find that when households and firms have expectations that are close to rational expectations, there is little difference between the behaviour of the 'Euler equation' and 'long-horizon' versions of the model. This means that the properties of key variables such as consumption and inflation — for example, in response to a change in the interest rate set by the monetary policy maker — are very similar, regardless of the assumptions we make about expectations. But when households and firms use expectations that are further away from rational expectations, the differences between the properties of the 'Euler equation' and 'long-horizon' versions become larger. This key result has implications for economic model builders. For cases in which households and firms have expectations of future income and costs that are very different from rational expectations of those variables, the model builder should choose the approach carefully.

Misperceptions, heterogeneous expectations and macroeconomic dynamics

Summary of Working Paper no. 449 Richard Harrison and Tim Taylor

An important question for economic policy makers is the extent to which the expectations of key decision makers in the economy affect — and are affected by — economic outturns. In particular, it is possible that mistaken beliefs about the behaviour of the economy can influence the behaviour of households and firms in a self-fulfilling manner. For example, a belief that inflation will be more persistent could influence price-setting behaviour so that actual inflation turns out to be more persistent. Such a feedback could reinforce the initial belief causing more households and firms to believe that inflation will be persistent.

This type of mechanism is illustrated in the following quote from the Bank of England's February 2008 *Inflation Report*: 'If households' and businesses' medium-term inflation expectations are heavily influenced by their recent experience, then repeated above-target outturns may cause them to place weight on the assumption that inflation will be persistently above [the inflation target of] 2%. If those expectations were built into higher wages and prices, that would raise medium-term inflationary pressures.'

To investigate this phenomenon, we build a small macroeconomic model in which the decisions of households and firms depend on their expectations for future income and costs, so that spending and price-setting decisions depend on expectations extending into the distant future. We assume that, to form their expectations, households and firms have access to a small set of alternative 'predictors'. These predictors are simple forecasting equations for relevant variables (for example, future inflation could be forecast by inputting recent observations for inflation into a simple equation). Households and firms choose between these predictors based on their recent forecasting performance. So a predictor that has forecast (say) inflation very well over the past few quarters will tend to be used more than a predictor with a worse forecasting record.

This 'dynamic predictor selection' creates the possibility of a feedback process between beliefs about the behaviour of the economy and its actual behaviour. We find that it is straightforward to generate this type of effect in our model under the assumption that households and firms choose between two predictors. The first predictor has very good properties when used by all households and firms. Its forecasting performance is close to the best possible predictor (the 'rational expectation'). The second predictor is a 'misperceptions predictor' which embodies a mistaken belief that inflation is more persistent. When we simulate the model, we are able to generate occasional periods of high, volatile and persistent inflation. This occurs when (random) shocks generate enough persistence in the inflation rate observed by households and firms to lead more of them to choose expectations based on the misperceptions predictor.

Forecasting UK GDP growth, inflation and interest rates under structural change: a comparison of models with time-varying parameters

Summary of Working Paper no. 450 Alina Barnett, Haroon Mumtaz and Konstantinos Theodoridis

In recent years, a number of papers have applied econometric models that allow for changes in model parameters. In general, this literature has examined and investigated how the properties of key macroeconomic variables have changed over the past three decades. So the underlying econometric models in these studies have therefore been used in a descriptive role.

The aim of this paper, instead, is to consider if these sophisticated models can offer gains in a forecasting context — specifically, GDP growth, CPI inflation and the short-term interest rate relative to simpler econometric models that assume fixed parameters. We consider 24 forecasting models that differ along two dimensions. First, they model the time-variation in parameters in different ways and allow for either gradual or abrupt shifts. Second, some of the models incorporate more economic information than others and include a larger number of explanatory variables in an efficient manner while still allowing for time-varying parameters.

We estimate these models at every quarter from 1976 Q1 to 2007 Q4. At each point in time we use the estimates of each model to forecast GDP, CPI inflation and the short-term interest rate. We then construct the average squared deviation of these forecasts from the observed value relative to forecasts from a simple benchmark model.

A comparison of this statistic across the 24 forecasting models indicates that allowing for time-varying parameters can lead to gains in forecasting. In particular, models that incorporate a gradual change in parameters and also include a large set of explanatory variables do particularly well as far as the inflation forecast is concerned, recording gains (over the benchmark) which are significant from a statistical point of view. Models that include this extra information also appear to be useful in forecasting interest rates. Models that incorporate more abrupt changes in parameters can do well when forecasting GDP growth. This feature also appears to surface during the financial crisis of 2008–09 when this type of parameter variation proves helpful in predicting the large contraction in GDP growth.

Neutral technology shocks and employment dynamics: results based on an RBC identification scheme

Summary of Working Paper no. 453 Haroon Mumtaz and Francesco Zanetti

Estimating the impact of changes in technology on the economy is one of the key aims of recent empirical research. And policymakers are equally interested, because in order to determine the appropriate stance of monetary policy it is essential to know what shocks are hitting the economy, and what their impact will be. The consensus from this literature is that the estimated impact can depend quite heavily on the way changes or shocks to technology are measured.

This paper contributes to this strand of the literature by proposing an improved procedure for measuring shocks to technology. In particular, we use information from a theoretical model of the business cycle which embeds labour market frictions to disentangle changes in technology from other shocks hitting the economy.

The estimation method comprises the following steps. First, we use the theoretical model characterised by search and matching frictions in the labour market to gauge the impact of the technology shock on vacancies, labour market tightness and other key macroeconomic aggregates. Second, we impose the predicted movements in these variables on US data, which has been the subject of many studies in the past. This is done via an empirical model referred to as a vector autoregression (VAR) where each included variable depends on the past values of all variables in the model. By using restrictions implied by economic theory, we can identify different types of shock, thus making the model a 'structural' VAR (an SVAR). The restrictions that we use are on the signs of impacts over particular time horizons. The SVAR is then used to estimate the response of key macroeconomic variables to technology shocks. The resulting responses of key macroeconomic variables provide us an approximation of the variables' responses to a change in technology in the United States.

Our main results are as follows. A positive shock to technology which affects labour productivity acts to increase GDP, investment, consumption and employment. This shock explains around 30% to 60% of the variation in each of these variables. This result is robust to a number of different configurations of the benchmark model and transformations of the data, such as controlling for long cycles in the data, choosing different time lags in the VAR, splitting the sample period, using alternative measures of labour market variables, and extending the length of sign restrictions on the SVAR.

One innovation is that we extend the benchmark model to allow the variance of the technology shock to change over time. We find that this shock played an important role in driving the volatility of US output during the 1970s and the 1980s. In particular, the volatility of technology declined since the early 1990s, which could explain the declined macroeconomic volatility over the same period, as highlighted in related studies.

Fixed interest rates over finite horizons

Summary of Working Paper no. 454 Andrew P Blake

Two natural questions to ask about monetary policy are 'what would happen to inflation if interest rates were a bit higher than forecast?' and 'what are the implications of interest rates not changing for some period of time?'. Satisfactory quantitative answers to both of these questions are, perhaps surprisingly, hard to come by. With many widely used forecasting models, this is not a problem. For example, the commonly used vector autoregression (VAR) — a system of equations explaining a set of interrelated variables — would allow us to simply impose a path for one of the variables with no practical consequences. But for policy we need to have a proper economic understanding, and one way of acquiring that is via a 'structural' model, which a VAR is not. Moreover, modern economics recognises the importance of forward-looking behaviour and expectations. Models where the forward-looking behaviour of agents helps explain the dynamic evolution of all variables in a coherent, equilibrium way are known as rational expectations (RE) models. Using an RE model to answer the questions just posed requires a forecaster to solve a number of quite difficult conceptual problems.

Using a general equilibrium RE model it is difficult to formalise how a higher (or indeed a fixed) interest rate is achieved. This is because such models are usually solved incorporating a monetary policy rule. These rules are conditional, and react to variables policymakers care about. Often, they are versions of the well-known Taylor rule that feeds back from inflation and growth. Departing from these rules to induce interest rates that are different from those already implied is hard to manage, and even if the technical problems are overcome it can be that the results sometimes seem perverse. Essentially, we cannot just 'fix' interest rates, as we can with VARs. In a structural model we have to have a coherent explanation of why interest rates follow the path they do (rather than what is implied by the policy rule embedded in the model). And the problem is compounded by the fact that behaviour in the model depends on what agents expect to happen after the fixed-rate path ends.

But the questions we began with are good ones that need reasonable answers. This paper explores a number of potential resolutions to modelling partially fixed interest rates in a common framework. These include imposing a sequence of anticipated or unanticipated interest rate 'shocks' that deliver the desired path, using a shock for each period the path is fixed, which seems a natural way to handle things. Unfortunately, when the strengths and weaknesses of different existing methods are compared they are all found wanting, either because they imply excessively volatile or counterintuitive forecasts. So a new approach is developed that restores more normal behaviour; but at the cost of introducing a new problem.

The new approach takes as a starting point that permanently fixed interest rates imply a well-defined trade-off between inflation and output growth, but do not imply any particular level of inflation. This is a well-known problem but (as we show) does not automatically apply in finite horizon problems, the case relevant for policymakers who publish fixed interest rate forecasts. Although at first sight the approach may seem somewhat perverse, the paper shows how to make sure it *does* apply for such problems. It can again be done by setting shocks, but using one more than the number of periods the rate is fixed; or by using a rule that specifically targets the interest rate, again for one period longer than the fixed-rate period. This restores intuitively sensible paths; but at the cost of introducing an equilibrium selection problem. This arises because when we use more shocks than we 'need' to fix rates, there are an infinity of well-behaved solutions that the forecaster must choose between. Equivalently, there are an infinity of rules we could use. A degree of arbitrariness in the selected solution is then inevitable. This is not as bad as it seems, though, as some paths are more 'sensible' than others (eg a path that is close to that implied by a Taylor rule). Nevertheless, the paper concludes that there is no easy solution to the finite horizon problem, and any answer to the questions we started with must inevitably be strongly caveated.