Asset pricing implications of a New Keynesian model

Summary of Working Paper no. 326 Bianca De Paoli, Alasdair Scott and Olaf Weeken

Macroeconomic models are widely used for policy advice. They are designed so that the behaviour of the model economy broadly matches that observed in economic data. But in many cases their implications for asset prices are not well understood. In particular, even though risk is an aspect of everyday life, these models tend to be silent about risk premia, ie the extra return investors require on risky assets, such as shares, to provide over and above the return obtainable from a riskless asset. Given the prominence of such models in policy advice, it is important that we develop a better understanding of their implications for asset prices. This paper investigates how asset prices are linked to the sources of economic uncertainty and the structure of the macroeconomy.

The model analysed in this paper is a typical macroeconomic model, a so-called New Keynesian model. It depicts optimising households and firms operating in goods and labour markets that exhibit some monopolistic behaviour. Also, in this framework, rigidities prevent real variables, such as consumption, labour and investment, and nominal variables, such as prices, from instantaneously adjusting to economic disturbances. In contrast to the optimising behaviour of households and firms, the central bank is assumed to follow a simple rule in which it adjusts a short-term interest rate to bring inflation back to target. The model represents a so-called closed economy in which households do not trade goods or assets with the outside world. Households can invest in a domestic equity index, nominal and real bonds of different maturities and a risk-free asset. Households use these financial assets to smooth consumption over time, selling assets to finance consumption when times are bad and

purchasing assets when times are good. There are two sources of uncertainty considered in the paper: a temporary increase or decrease in productivity and a temporary deviation by the central bank from its usual behaviour. Contrary to many works in the literature, the model is solved in a way that takes account of the effects of uncertainty on the economy, thus capturing the different risk premia associated with the assets under consideration. This implies that the size and sign of these risk premia depend on how well an asset helps households to smooth consumption and the quantity of risk present in the economy. Assets that are expected to pay well in bad times when growth is expected to be low are more highly valued than assets that are expected to pay well in good times. This is the familiar result that risk premia depend on the comovements between economic variables and asset returns.

The paper demonstrates how risk premia are linked to the two sources of uncertainty and the rigidities in the model. In particular, the paper highlights that because different economic shocks imply different comovements between asset returns and growth, the source of shocks will be an important determinant of risk premia. It also demonstrates how the size of these premia depend on how the economy deals with uncertainty, which in turn depends on the form of economic frictions and rigidities present in the economy. For example, real rigidities that prevent goods and labour markets adjusting after shocks will increase risk premia regardless of the source of the shocks. On the other hand, nominal rigidities, that slow down price adjustment to shocks, increases risk premia when the economy is hit by productivity shocks, but reduces risk premia in the presence of monetary policy shocks.

Summary of Working Paper no. 327 Lavan Mahadeva

The bulk of evidence suggests that some combination of improvements in the monetary policy making institutions (more transparency, enhanced credibility and stronger accountability), consolidation in fiscal policy, a more benign international economic environment and technological improvements in the dissemination of data releases has reduced interest rate volatility in the United Kingdom in the 1990s. But it is difficult to isolate the role of transparency by itself in reducing interest rate volatility. Still, arguably, the incidence of market surprises is greater than one might expect given the extent of these improvements. Might this be because transparency affects surprises differently when it has improved from low levels than when further improvements are made at high levels of transparency? To help answer this question, this paper presents a theory to link improvements in transparency to how well financial markets predict policy rates.

In the paper, policymakers may have a different interpretation to private agents as to what an economic shock means for interest rates, even though data is commonly available. There is then a role for policy announcements (both the interest rate decision and the surrounding explanations) to act as a beacon communicating information to agents about the policymakers' preferences. Policymakers take account of this feature when they determine first, how to set interest rates and second, in choosing how transparent they want to be in explaining that decision. What determines how transparent policymakers want to be with their explanations? The paper shows that policymakers are more likely to be transparent in their explanations if they follow a transparent objective. The paper describes an improvement from a bad regime where policymakers are allowed to follow their own secret, unpredictable inflation objective to a good regime where the inflation target is publicly known and fixed. Policymakers in bad regimes will be inclined to also be less transparent in explaining their actions while policymakers in good regimes have a strong incentive to be transparent about their explanations.

Improving transparency in objectives typically lowers the volatility of market interest rates and implies less market surprises. But the paper also shows that this happens at a decreasing rate; improvements in transparency from already high levels have less effect on reducing the likelihood of market surprises than when transparency is improved from very low levels. At high levels of transparency, agents rely more on the cleaner signal and this results in greater sensitivity of agents' expectations. This feeds back onto interest rates offsetting the effects of greater transparency. In general, though, improving transparency leads to more precisely formed inflation expectations. In this sense greater transparency is good for welfare.

Cash-in-the-market pricing and optimal resolution of bank failures

Summary of Working Paper no. 328 Viral Acharya and Tanju Yorulmazer

The idea that rescuing troubled banks can create incentives for excessive risk-taking is widely spread. However, empirical evidence suggests that regulatory actions taken in response to banking problems vary significantly. In many episodes, regulatory actions appear to depend on whether the problems arise from idiosyncratic reasons specific to particular institutions or from aggregate reasons with potential threats to the whole system. When faced with individual bank failures, authorities usually seek a private sector resolution, whereas government involvement is an important feature of the resolution process during financial crises that affect a significant portion of the banking industry, that is, during crises that are systemic in nature. We argue in this paper that this difference in regulatory actions arises from the fact that resolution options open for an isolated failure of a single institution are different from those available when facing a systemic failure. When only a few banks fail, these banks can be acquired by the surviving banks. However, when the crisis is systemic, that is, for a large number of failures, the liquidity of surviving banks may not be enough for them to acquire all failed banks at the full price. This may lead to the price of failed banks' assets being determined by the available liquidity in the market, resulting in 'cash-in-the-market' pricing of failed banks' assets. Furthermore, during systemic crises, it is more likely that investors outside the banking sector, who are liquidity endowed but potentially not the most efficient users of these assets, end up purchasing some failed banks' assets, leading to social welfare losses associated with the misallocation of banking assets.

Thus, when the banking crisis is systemic in nature, there are 'too many (banks) to liquidate' and bailing out some of the failed banks by the authorities may be optimal in order to avoid allocation inefficiencies. However, this bailout policy may be suboptimal, because it may induce banks to herd by lending to similar industries or betting on common risks such as interest and mortgage rates, in order to increase the likelihood of being bailed out. This in turn increases the likelihood of experiencing systemic banking crises in the first place. We show in this paper that other regulatory options such as the provision of liquidity to surviving banks to be used in acquiring failed banks' assets can mitigate this problem. We show that this policy is equivalent to the bailout policy, but gives banks incentives to differentiate, rather than to herd.

In this paper, we formalise these ideas in a framework wherein the optimal bank failure resolution policies and the cash-in-the-market pricing are endogenously derived. We consider a variety of resolution policies that broadly cover the entire spectrum of policies employed by regulators to resolve bank failures. In particular, failed banks may be closed, in which case their assets are sold to surviving banks and outsiders at market-clearing prices, or failed banks may be bailed out, in which case their owners are allowed to continue operating the banks. The regulator may also provide liquidity to surviving banks to be used in acquiring failed banks' assets. We show that by virtue of assisting surviving banks in acquiring more failed banks, the liquidity provision policy increases the anticipated surplus for banks in states with cash-in-the-market prices. In turn, this mitigates banks' incentives to herd.

The impact of yuan revaluation on the Asian region

Summary of Working Paper no. 329 Glenn Hoggarth and Hui Tong

This paper aims to analyse how an appreciation of the yuan affects the exports of other Asian countries, by controlling for the rapid change in the structure of China's exports.

China has increasingly been acting as a 'world factory' since the early 1990s. Because of this, if China's exports fall following a yuan appreciation, then its demand for upstream intermediate and capital goods may decline as well, even if these imported goods become less expensive. Therefore, it is possible that yuan appreciation will have a much smaller positive impact (or perhaps even a negative one) on the trade surplus of high-income Asian capital-goods exporters than on low-income Asian consumer-goods exporters — a hypothesis that this paper will examine.

There is extensive empirical research on how exchange rate movements affect the trade balance in general and those of Asia and China in particular. But the literature so far has shared a limitation: the estimation tends to be based on a relatively long historical period (25 years or more). But there has been a significant change in the structure of international trade, and of Asian trade in particular, over this period. This paper aims to fill this gap by using a panel estimation that uses a large data sample that controls for the change in the commodity structure of trade.

Bilateral export and import equations are estimated between China and nine Asian countries: India, Indonesia, Japan, Korea, Malaysia, Pakistan, the Philippines, Singapore, and Thailand since the early 1990s. Competition between China and these countries in third markets is also estimated.

Three related empirical models are examined: China's own exports, China's imports, and the competition between Asian countries and China in third markets. The results are consistent with the supply-chain story, whereby China imports capital goods from advanced Asian countries to facilitate the production of consumer goods exported to third markets. When the yuan appreciates, China's exports fall, which then reduces China's demand for upstream capital goods. Consequently, exporters of mainly capital goods to China, such as Japan and Korea, are found to be adversely affected by a yuan appreciation. There is also little evidence that Asian countries benefit from yuan's appreciation in their exports to third markets.

Escaping Nash and volatile inflation

Summary of Working Paper no. 330 Martin Ellison and Tony Yates

This paper builds on a story Tom Sargent told in his book The conquest of American inflation. That book seeks to explain the rise and fall of inflation in the United States and provides a cautionary tale for those who are confident that inflation will not rise again in the future. The story involves a monetary authority that sets policy believing, incorrectly, that higher inflation can buy permanently lower unemployment. The monetary authority takes a model like this to the data and updates its estimates of this long-run trade-off. In the model, inflation is high when high inflation appears to buy lower unemployment: when the benefits of high inflation appear greatest. Inflation is low when it appears to the central bank that unemployment is unrelated to inflation. Unbeknownst to the policymaker in this model, only surprises in inflation differences between what was expected by the private sector in the model, and the outturn — affect unemployment. Periodically, unobserved shocks come along that offset 'mistakes' that the central bank makes setting inflation, and this makes it look to the central bank as though inflation has no effect on unemployment. So the central bank chooses low inflation.

Our modification to this story is to adapt Sargent's model to explain why the volatility of inflation seems to be high when the level of inflation is high, and *vice versa*, a matter on which the original model is necessarily silent. We add a shock to the model which the central bank sees, but to which the private sector cannot react, and which the central bank seeks to offset using its inflation policy. This shock brings the model a little closer to reality, since it is likely that monetary policy has in the past sought to help stabilise inflation and the macroeconomy. In our model, at times when the central bank thinks the benefits to high average inflation are greatest, it also thinks it can use inflation to stabilise unemployment more effectively, and therefore chooses not only high inflation, but also volatile inflation. Conversely, when it appears to the monetary authority that high inflation does not buy any reduction in long-run unemployment (when the trade-off appears to be vertical), it chooses not only low inflation, as in the original Sargent model, but also more stable inflation, since it sees no point in trying to use inflation to offset the shock to unemployment.

We also document that escapes from high and volatile inflation to low and stable inflation in our model are less likely if the variance of the observed shock to unemployment is high. A rough intuition for this effect is that the more variable is the observed shock to unemployment, the more variable the central bank chooses inflation to be. This generates larger inflation surprises, and those surprises translate into larger movements in unemployment. What the central bank sees is a strong correlation between inflation itself and unemployment, making the long-run trade-off look promising for the central bank, and that reinforces its belief in the benefits of high and volatile inflation (which it incorrectly thinks will translate into low unemployment). The only thing that can undo this correlation in the model is the unobserved shock to unemployment. If this shock is large enough, then from time to time, it can offset the effect of an inflation surprise on unemployment and eliminate the correlation the central bank sees between inflation and unemployment. But the more variable is the observed shock relative to the unobserved shock, the less powerful the latter is in wiping out this correlation. If we were to take this model to the real world, then this feature would suggest that part of the reason inflation became lower and more stable was because the volatility of observed shocks to the supply side fell.

Taken literally, the model has the rather gloomy prediction for the future that there will forever be bouts of high and volatile inflation. But what our model misses out, like the Sargent model it derives from, is the revolution in the institutional design and economic understanding underpinning monetary policy that took place in recent decades.

Wage flexibility in Britain: some micro and macro evidence

Summary of Working Paper no. 331 Mark E Schweitzer

Increased wage flexibility was often cited as the main reason behind weaker inflationary pressures in the 1990s. Wage flexibility can be defined as either a micro or macroeconomic concept; each is quite distinct, although potentially related. For example, in a Phillips curve an increase in macroeconomic wage flexibility is often captured by a larger wage response for a given unemployment rate. On the other hand, increased microeconomic wage flexibility is usually identified by the lack of evidence for wage rigidities, such as limited evidence of a spike in the distribution of wage changes at zero. An abundance of zero wage changes in the data might indicate an inability to adjust wages in a timely manner.

This paper uses data from the British New Earnings Survey from 1975–2000 to derive both macro and micro measures of wage rigidities. Because the data span a 26-year period, the behaviour of micro and macro flexibility measures can be compared over time. In addition, we can investigate whether there is any evidence that the behaviour of either measure of flexibility has shifted over time. To keep the analysis simple, we consider whether there has been a one-off shift in wage flexibility. Regional wage rates do not appear to have been more responsive to regional unemployment levels in the 1990s than in earlier years. Instead, estimates focusing on the 1990s reveal no statistically significant aggregate wage response to regional unemployment levels. The overarching conclusion is that macroeconomic tests leave much to explain, but the estimates have revealed some patterns that are worth trying to reconcile with other sources of evidence.

In individual-level wage data, there is evidence of so-called nominal wage rigidity. In contrast to previous findings, the evidence is generally stronger for the latter half of the sample period. The degree of nominal wage rigidity is somewhat smaller than found in related research using data for the United States over the full sample period, but the estimates over the recent past are similar. The evidence of nominal rigidities is somewhat surprising in the light of the popular view that the UK labour market has become gradually more flexible.