Elasticities, markups and technical progress: evidence from a state-space approach

Summary of Working Paper no. 300 Colin Ellis

UK monetary policy is concerned with keeping inflation — the rate of increase in prices — on target at 2% a year. So it is important for policymakers to consider how firms set prices. Typically, economists work with models that assume companies set their output prices as a markup over marginal cost — that is, the cost of producing an extra unit of output.

In most economic models, that markup is assumed to be fixed, at least on average over a long period of time. But in practice, it is possible that the markup could have changed over time, for example if competition between companies becomes more intense. At the same time, standard economic models often impose an assumption about production technology: in particular, how easy companies find it to swap between machines (capital) and workers (labour) when they produce their output. This is called the elasticity of substitution in production. In fact, any assumption about the markup will affect the estimated elasticity of substitution in a model, and *vice versa*.

This paper proposes a new approach, where the markup and elasticity are jointly estimated. In particular, the markup is allowed to (potentially) vary over the past 30 years. The model is estimated using so-called 'state-space' techniques, which allow the unobserved markup to be modelled using UK data on prices, wages and other macroeconomic variables. The estimation results are very different from what standard approaches find — in particular, the state-space approach suggests that the aggregate markup in the UK economy has fallen by around a quarter since the early 1970s, and that firms find it harder to swap between capital and labour than is often assumed. In addition, the model also lets technical progress in the economy — a gauge of the efficiency with which firms use capital and labour to make output — be estimated in a more realistic manner than in most models. This turns out to be crucial — the usual approach in other work, of simply including a time trend in the model, is shown to give misleading results.

The key results from using the state-space model are robust to a number of consistency checks, such as the degree of tightness in the labour market, looking at the private sector rather than the economy as a whole, and measuring how useful machines are in production, rather than what they are worth. Given that the model focuses on long-run effects, data from the 19th century are used to check that running the model from 1970 is not misleading. Finally, the model is applied to US data, again retrieving plausible results.

This new approach of treating the markup as unobserved and estimating it at the same time as production technology yields several insights. First, the markup in the United Kingdom has fallen over the past 30 years or so. This implies that the unit labour cost of production — essentially the pay workers receive for each unit of output they produce — has not always been a good guide to the marginal cost of production, despite it being widely used to proxy marginal cost in previous work. Second, firms find it harder to swap between capital and labour in production than most other estimates suggest. Finally, using a time trend to proxy technical progress can be very misleading.

The welfare benefits of stable and efficient payment systems

Summary of Working Paper no. 301 Stephen Millard and Matthew Willison

The Bank of England's second core purpose is to maintain the stability of the financial system, both domestic and international. A key aspect of financial stability is the ability of consumers, firms and the government to continue making payments to each other in the presence of shocks both external to and emanating from within the systems through which such payments are made. Examples of such shocks could include bankruptcy of payment system participants; liquidity shortages among participants or problems with their operations; events in the wider economy that lead to changes in the profitability or liquidity holdings of participants in the system; or operational problems within the system leading to its temporary closure.

In this paper, we introduce a payment system into a recently developed theoretical model of banks and examine the ability of agents in an economy to make payments to each other in the presence of operational problems within this payment system. In the model agents have a choice between two means of making payments — cash and an alternative — but only one, cash, can be stolen. The safe alternative to cash is referred to as 'cheques' but, in essence, this can be thought of as any reliable interbank payment system. The introduction of a payment system (and banks) enables agents to more easily make payments between each other.

But the payment system in their model is risk free. In the real world this is not the case. In particular, such systems can suffer from operational problems, the focus of this paper. There is a risk that the payment systems temporarily fail to function for some reason and payments cannot be made. In this paper, we model the possibility of shocks to the payment system as a probability that the payment system fails to function when a buyer and producer who meet and agree to trade would like it to. Agents do not know whether the system will function or not when they choose whether to use it rather than using cash. We show that agents have an incentive to use the payment system if it is sufficiently cheap to use and/or sufficiently reliable. We also derive lower bounds on the probability that the payment system functions (given the cost of using it) that are consistent with buyers choosing to use it.

Finally, we compare social welfare with and without the payment system. The presence of a safe and reliable system for transferring money can make people prepared to hold and use money in situations where the presence of thieves would have otherwise stopped this from happening. In such cases, the presence of a payment system unambiguously increases social welfare since it expands the number of trades occurring in the economy. We find that the more reliable the system, the more likely this is to happen. Using our model, we then calculate the welfare gains resulting from an increase in stability. When money is accepted as a medium of exchange in the absence of a payment system, social welfare can increase or decrease with the introduction of a payment system. In this case, the addition of a payment system will not expand the number of trades that occur in our model; so there will be no social benefit arising from this channel. Social welfare will only increase if the reduction in deadweight loss caused by theft in the economy (a cost that thieves incur when they steal successfully) is sufficiently greater than the costs of using the payment system (including both the direct costs of using the system and costs related to system failures). Again, we show that this is more likely for a more reliable system, and calculate how welfare increases as stability is increased.

International and intranational consumption risk sharing: the evidence for the United Kingdom and OECD

Summary of Working Paper no. 302 Vincent Labhard and Michael Sawicki

'Consumption risk sharing' refers to the ability of households to protect their consumption against shocks to their income. This could take the form of holding equity claims on output that is unrelated to their income, by receiving transfers from other agents or by borrowing and lending. Successful risk sharing should imply a smooth pattern of consumption, not greatly affected by fluctuations in households' incomes. However, empirical work has shown that this does not appear to be the case. For example, households hold a much smaller proportion of foreign equity than would be expected if they decided on their asset holdings on the basis of risk and return, a phenomenon known in the literature as the home bias in equities. A related puzzle is the home bias in consumption. Under full risk sharing, domestic consumption growth should be more highly correlated with foreign consumption growth than with domestic output growth, but the empirical evidence again suggests the contrary: cross-country consumption growth correlations are relatively low, and often lower than correlations with domestic output growth. This would suggest that 'idiosyncratic' — ie country or region-specific — output shocks are not effectively smoothed away and hence materially affect consumption.

Measuring the overall extent of risk sharing in consumption is interesting for monetary policy makers because of its impact on the transmission of shocks. Gauging the extent of risk sharing — as well as the channels through which it occurs — can help us understand business cycle developments and imbalances in an economy by informing us about how consumption is likely to respond to country or region-specific output shocks, and can shed light on how much policymakers might need to react to such shocks. As another potential adjustment mechanism against country-specific shocks, a large enough degree of international risk sharing may mitigate the effects of such shocks if other adjustment mechanisms, for example exchange rate flexibility or labour mobility, are absent or limited. And finally, understanding the channels by which risk sharing is achieved would also provide an insight into the effects of capital and credit market integration, both domestically and internationally.

In this paper, we present the empirical evidence for consumption risk sharing by UK consumers, both between the United Kingdom and other countries (internationally) and across regions of the United Kingdom (domestically). Such evidence can tell us whether risk sharing acts as an absorber of country or region-specific shocks in the United Kingdom. The key questions we seek to answer are whether there is more or less risk sharing domestically than internationally; through which channels it occurs; whether risk sharing has increased or decreased over time; and finally, whether these estimation results are robust. We address these questions by employing two specific methodologies. First, we use an established panel regression analysis — using a data set spanning a set of countries or regions across time — to illustrate both the extent of risk sharing and the channels through which it is carried out, updating existing UK results with more recent data. The incremental information on the channels by which risk sharing occurs means that we prefer this methodology to simple correlations of consumption and output data. However, as with correlations, the panel analysis may be distorted by factors affecting output or consumption but not related to risk sharing, such as changes in household preferences and measurement error in the data. This prompts us to utilise a second, more recent, methodology, which takes these influences explicitly into account using a factor model — a technique which aims to separate out the key drivers or factors (here at the regional and national level, for example) from a potentially large set of data — on consumption and output. Applying this factor model to the United Kingdom and OECD data is the main contribution of this paper.

We find that there is more risk sharing across the UK regions than between the United Kingdom and OECD countries. This baseline result is robust to accounting for the possible impact of measurement error and changes in household preferences, and consistent with results reported in the previous work. We find that the main mechanism of regional risk sharing operates via cross-regional asset holdings. Internationally, the main source of income smoothing comes from international borrowing and lending. Consistent with previous work in the field, we find tentative evidence that risk sharing has declined over time, although the importance of capital market smoothing has gradually increased, consistent with recent increases in capital market integration. However, these trends may require caution in interpretation, because the methodologies we use may not fully detect changes in the nature of the risks to output occurring during the course of our sample.

Finally, our paper also makes a separate contribution to the literature by illustrating the role of the choice of price measures (deflators) in estimating the true extent of risk sharing for the United Kingdom and OECD. While estimates of the extent of risk sharing within the United Kingdom are relatively invariant to the choice of deflator, using our preferred choice of deflator for the OECD data set yields higher estimates of risk sharing than typically reported in previous work.

Therefore, it appears as if in the United Kingdom regional consumption fluctuations may be largely unaffected by regional output fluctuations, and that UK consumption — while affected by global output fluctuations — may also be more robust than suggested in previous work.

The danger of inflating expectations of macroeconomic stability: heuristic switching in an overlapping generations monetary model

Summary of Working Paper no. 303 Alex Brazier, Richard Harrison, Mervyn King and Tony Yates

In the past decade we have witnessed a step change in macroeconomic performance. Both output growth and inflation have been much more stable than they were in the 1970s and 1980s. Policymakers — keen that this development should be durable — have tried to understand the causes of this 'Great Stability'.

Research has tended to place explanations into two groups: good luck in the form of fewer, smaller shocks; or better stewardship of the economy by governments and central banks. But as Ben Bernanke, Chairman of the Federal Reserve, has pointed out, apparently smaller shocks might actually be the result of better stewardship of the economy through anchoring of expectations about the future. With clear objectives for monetary and fiscal policy, expectations of the future need not be guided by what happens to the economy today. They can instead be guided by those clear objectives. When decisions made today by businesses and households depend on their expectations of the future, their actions will be less sensitive to past developments.

So thinking about the way expectations are formed is likely to be an important step towards understanding the 'Great Stability'. Standard economic models tend to assume that people form expectations using detailed knowledge about the way the economy works and the shocks that hit it. What if, instead, they use simple rules of thumb — or 'heuristics' — to form their expectations?

This paper explores that question using a very simple model in which output and inflation today depend on expectations of inflation tomorrow. Although the model is too abstract to explore questions about current monetary policy, it allows us to explore the more general issue of the role that expectations play in shaping economic performance.

In this model, predictions about future inflation affect decisions about what to do today. So if, other things being equal, inflation is expected to be high tomorrow, it will start to pick up today. Predicting future inflation is difficult because the economy is subject to temporary, unpredictable shocks. So people are assumed to use one of two heuristics. The first is to assume that inflation tomorrow will be the same as inflation yesterday — the 'lagged inflation' heuristic. The second is to assume that inflation tomorrow will equal a target announced by the central bank — the 'inflation target' heuristic.

It seems reasonable to assume that people will only adopt a particular heuristic if it would have predicted inflation well over the past. That opens up the possibility that people might switch between them. Sometimes, when inflation has been close to target, they are likely to use the inflation target heuristic. When they do, we know that it helps to keep inflation stable because it acts as an anchor for expectations. Sometimes, however, shocks will move inflation away from the target. If people then switch to using a lagged inflation heuristic, there will not be a firm anchor for expectations. This means that there are periods like 'Great Stabilities' in which inflation is very stable but these are interspersed with periods of greater volatility.

In our experiments, the inflation target is only used as a heuristic if it would have performed better than others in the past. There is no guarantee that it will be used. But the announcement of an inflation target at least opens up the possibility of more stable periods of economic performance. And, as such, inflation in an economy with an inflation target tends to be more stable than in an economy without a target.

That illustrates the importance of the monetary policy framework in this model. Given that framework, what does this simple model say about how a central bank should operate? It is not possible for us to draw conclusions for the conduct of monetary policy in the real world because the model is so abstract. But we illustrate how, in this model, monetary policy can better stabilise the economy by responding to inflation expectations. This contrasts with many standard economic models, in which inflation expectations contain no information about the state of the economy that is not already apparent in other indicators. But in our model, inflation expectations do contain information about the state of the economy. They indicate which heuristic people are using and, therefore, how the economy will respond to shocks.

Procyclicality, collateral values and financial stability

Summary of Working Paper no. 304 Prasanna Gai, Peter Kondor and Nicholas Vause

A financial system plays a highly beneficial role in an economy by helping to transfer resources to sectors where they can be used most productively, with transfers taking place both across time and potential states of the world that could materialise. In principle, a perfect financial system could insure the constituent sectors of an economy from the idiosyncratic risks that they face, so that fluctuations in economic activity at the macroeconomic level would reflect only systematic shocks, ie those that affect all sectors. But financial systems operate under frictions such as asymmetric information, where some market participants are better informed than others, and this makes financial contracts costly to monitor and enforce. A practical view that appears to be becoming more widespread suggests that when financial systems operate with frictions, economic shocks can be amplified and propagated, exaggerating economic upturns and prolonging the severity of economic downturns, and leave economies more vulnerable to such shocks during expansionary phases of the business cycle.

This paper outlines a model that analyses both how macroeconomic shocks can be amplified and how procyclical macroeconomic risk can be generated within a macro-financial system. The model is constructed so that shocks that boost the productivity of one sector adversely affect the productivity of the other sectors. Thus, a series of shocks that raise the output of one sector, such as a clustering of technological innovations, will cause the economy to grow as this sector accounts for a greater and greater share of the total economy. And, as the economy becomes more concentrated, it becomes more vulnerable to the dominant sector being hit by an adverse shock at some point in the future.

The financial system in the model allows risk-averse entrepreneurs in the economy to insulate their balance sheets against uncertainty. But financial contracts must be supported by collateral, such as real estate, to ensure that promises to make payments in the future are credible. If the collateral asset is also used in production, a feedback loop between aggregate output and the value of collateral emerges. A key contribution of the paper is to show how such feedback loops are maintained in the presence of insurance markets. An initial decline in aggregate output reduces entrepreneurs' net worth and, hence, the price of the collateral asset, as demand for the asset for use in future production declines. The decline in the value of the collateral asset implies that producers are unable to obtain sufficient insurance, exposing balance sheets to shocks. Since entrepreneurs are risk-averse, their response to additional balance-sheet uncertainty is to reduce the scale of production. This leads to subsequent declines in the price of the collateral asset, completing the feedback loop. Any decline in its value as a result of incomplete insurance by one sector leads to inadequate insurance by other sectors. This externality increases the level of systemic risk across the economy. Systemic risk imposes welfare costs on the economy as it leads to inefficient production and results in balance-sheet uncertainty. Both aspects are captured by the model.

Bank capital, asset prices and monetary policy

Summary of Working Paper no. 305 David Aikman and Matthias Paustian

Do weak banks affect the transmission mechanism of monetary policy? Does bank lending merely reflect general macroeconomic conditions, or are there important feedback effects from banks to other macro variables? More generally, how should financial sector conditions influence the conduct of monetary policy? These questions are of long-standing interest to policymakers and they form the motivation for this paper. In order to study them, we develop a framework that explicitly models the role of banks in intermediating credit flows, and takes into account some possible frictions that are likely to exist between depositors, banks and borrowers.

In our model, the amount of capital held by banks and the creditworthiness of borrowers are both important ingredients in transmitting shocks throughout the economy. To see why, suppose that an unanticipated tightening of monetary policy (or some other adverse shock) leads to a decline in output, which then lowers the profitability of firms, triggers a fall in asset prices, and causes loan losses for banks. The accompanying reductions in borrower net worth and bank capital will have two effects. First, banks will be less willing to lend to borrowers whose creditworthiness has declined. And second, depositors will view banks as riskier institutions, and will readjust their portfolios out of bank deposits. We show by simulation that these effects are able to generate a significant second-round cutback in the flow of lending which exacerbates the initial downturn.

Intuitively, we might expect there to be a role for monetary policy in mitigating the second-round effects generated by these frictions. For instance, by aggressively cutting interest rates in a downturn, the central bank might be able to check the falls in asset prices and net worth associated with the shock, thereby partly cushioning the impact on aggregate demand. The cost of acting in this way, however, is higher inflation — at least in the short term. A key question for policymakers is therefore: how much of an increase in inflation volatility should be tolerated in order to reduce the volatility of output growth in this way?

The chief contribution of this paper is to tackle this question. We proceed in two steps. First, we assess the performance of monetary policy strategies that respond in a mechanical way to 'financial' variables such as asset prices or credit flows over-and-above consumer price inflation. It turns out that these simple monetary policy rules perform poorly if the goal of policy is maximising the wellbeing of economic agents in the model. Second, we use numerical techniques to analyse the properties of the 'optimal' monetary policy implied by the model. Our main finding is that a central bank acting in this optimal way will tolerate only a very small amount of inflation volatility. Furthermore, the 'trade-off' implied by our model is very steep in the sense that the reduction in output growth volatility achieved by allowing inflation to become more volatile is very small. Given that similar results have been reported for models that abstract from banks — and in fact credit market imperfections altogether — we conclude that assigning a non-trivial role for these frictions need not materially affect optimal monetary policy. This suggests that policies that work well in 'normal' times are likely to continue working well in a situation where weak banks are limiting the expansion of credit.

Consumption excess sensitivity, liquidity constraints and the collateral role of housing

Summary of Working Paper no. 306 Andrew Benito and Haroon Mumtaz

Consumer expenditure is the dominant component of aggregate demand, and as such, understanding consumption plays a central role in understanding the behaviour of the macroeconomy. That requires a good understanding of how households form their consumption plans. The most influential way of thinking about how households do that is through the life-cycle model. That is based on the idea that households are forward looking and wish to avoid changes in the satisfaction they get from consumption during their lives. In that way, households smooth their consumption.

It has long been recognised that some households may not smooth their consumption to the full extent implied by the life-cycle model. The first aim of this paper is to estimate what percentage of households in the United Kingdom do not smooth their consumption in such a manner. Among other things, that is important for understanding how households will adjust their spending in reaction to shocks that affect their income.

In recent years there has been increasing interest in the role of housing and its relationship with consumption. On several occasions in the past, consumption and house prices have moved together. But Monetary Policy Committee discussions have noted that the reduced-form relationship between consumption and house prices has recently appeared weaker than in earlier periods.

There are various channels through which house prices can influence consumption, notably the so-called collateral channel, and common determinants of both housing demand and consumption. One view has it that house prices are an asset price for an essential commodity, shelter, and that they merely reflect macroeconomic conditions with no special role of their own. But on another view, there is an important causal effect of housing in providing collateral. That allows credit to be obtained on more favourable terms and supports consumption. That role may be particularly strong, or only exist at all, for those that might otherwise have been constrained by the availability of credit. Among other things, this collateral channel could amplify the effects of monetary policy on the economy. However, there is little evidence on whether housing equity fulfils this role and how it affects households' consumption plans. A further aim of this paper is to use microdata to confront the implication of the collateral hypothesis that housing capital gains should affect those that are liquidity constrained differently from those that are not liquidity constrained.

If households smooth consumption as the life-cycle model implies, then current consumption plans should not react to past news about income: that should already be incorporated into households' consumption plans. In this paper, we explicitly model the likelihood that a household's behaviour falls into one of two 'regimes' according to whether the household displays 'excess sensitivity' to recent income news or not (ie, whether it fails to smooth consumption). We find that around 20%–40% of households display excess sensitivity. These households are liquidity constrained or saving for other precautionary reasons. The former are households who would like to borrow to smooth consumption but cannot, or face a relatively high interest rate which puts them off borrowing. The latter are those who are reluctant to borrow because of the risks of large amounts of debt when future income or expenses are uncertain. They can be said to have a 'self-imposed' liquidity constraint and instead want to accumulate their buffer of assets. We find that households are more likely to fall into either group if they are without liquid assets, have negative home equity, are young, unmarried, non-white or are degree-educated.

Regarding the collateral channel, in addition to the effect of negative home equity in influencing the likelihood of being liquidity constrained, we also find evidence for the effect referred to above, that housing capital gains affect the consumption of those that are more likely to be liquidity constrained. That is direct evidence in support of the existence of a collateral channel.

Fiscal rules for debt sustainability in emerging markets: the impact of volatility and default risk

Summary of Working Paper no. 307 Adrian Penalver and Gregory Thwaites

The prospects of receiving full payment of emerging market sovereign debt cannot be established with certainty. In emerging market economies (EMEs), key macroeconomic variables — the primary budget balance, economic growth, inflation, domestic and foreign interest rates and the exchange rate — are typically volatile, making it difficult to predict the future with confidence. These macroeconomic variables are also usually correlated. For example, an adverse terms-of-trade shock can slow output growth, result in an exchange rate depreciation and raise the risk premium on interest rates, all of which will worsen a sovereign's debt position. The volatility, correlation and persistence of shocks in emerging markets mean that assessing debt sustainability on a single future path of these variables is too simplistic. Forecasts based solely on historical averages of these variables may therefore erroneously neglect a chance that sovereign debt and fiscal policy are unsustainable.

Using a simple econometric model estimated on a representative EME, this paper measures how uncertainty about the future and the effect of the risk of sovereign default on interest rates alters the probability of future debt to GDP outcomes. Simulations of this model under alternative fiscal policy regimes show that any stabilising fiscal policy must react strongly to innovations in the debt-GDP ratio. Forecast uncertainty and feedback from the debt level to real interest rates impose material constraints on the set of fiscal policy rules which stabilise debt.

These techniques and analysis have practical policy implications. A quantitative analysis of the uncertainty that surrounds debt projections could help the IMF when assessing members' debt sustainability before agreeing to financial assistance programmes. It could also support IMF surveillance of fiscal policy and thereby contribute to crisis prevention. The technique may also be useful to policymakers in EMEs when determining their medium-term fiscal policy strategy. It would be particularly useful if a country is considering the introduction of fiscal policy rules.

Optimal emerging market fiscal policy when trend output growth is unobserved

Summary of Working Paper no. 308 Gregory Thwaites

This paper is concerned with how fiscal policy in emerging markets should respond to changes in economic conditions. We model the behaviour of a fiscal authority in an emerging market economy (EME) who can borrow from other countries to smooth the effects of unexpected changes in residents' spending. We focus on the policy implications of (1) Aguiar and Gopinath's finding that GDP in emerging markets is characterised by relatively large and persistent shocks to the trend growth rate, and (2) that policymakers cannot directly observe the output gap or the trend GDP growth rate.

We have two key findings. First, we find that risk-averse policymakers who face EME-style output processes (ie changes in output are dominated by shocks to trend growth) should run tighter fiscal policies, with lower average debt-GDP ratios, than those in industrialised countries. This result is robust to agents' risk-averseness and dislike of holding debt, as well to the amount of real-time information that the policymaker has on what determines a change in output. In all cases, and particularly when the interest rate is very sensitive to the debt-GDP ratio, we find that the primary fiscal balance (ie excluding net debt interest payments) should respond strongly as the debt-GDP ratio moves above its long-run average. We find that the introduction of moderate shocks to the gross return on debt (eg those due to real-exchange rate shocks if debt is denominated in foreign currency) has little effect on the relationship between the primary fiscal balance and the debt-GDP ratio, unless the shocks happen at the same time as offsetting changes to the trend growth rate.

Second, in our baseline 'EME' model, we find that the primary fiscal balance of an optimising policymaker is countercyclical, despite changes in output being driven by shocks to trend growth. This appears to be true irrespective of the amount of information the policymaker has about trend growth and the output gap. This result contrasts with other papers, which have used optimising frameworks and the features of EME output processes to rationalise the observed procyclicality of EME fiscal policies or external balances. The result is somewhat sensitive to assumptions about debt intolerance and risk aversion; in particular, greater debt intolerance makes policy more countercyclical.

Our simulations also suggest that the welfare costs of naively running a fiscal policy that would be appropriate for an industrialised country are around 1% of average consumption. But this result is sensitive to assumptions about capital markets and risk aversion. We find that a simple rule-of-thumb policy that stabilises the debt-GDP ratio in every period results in smaller welfare losses than if the 'industrialised' policy is implemented.