Fundamental inflation uncertainty

Summary of Working Paper no. 309 Charlotta Groth, Jarkko Jääskelä and Paolo Surico

The Phillips curve, which relates inflation to some measure of real activity, plays an important role in modern economic theory. The relationship is also important for policymakers, as it serves as a useful description of the short-run inflation dynamics. In modern New Keynesian models, the Phillips curve is explicitly derived from the pricing decisions of firms, that set their prices as a mark-up over costs. At the aggregate level, the baseline New Keynesian Phillips Curve relates current inflation to lagged and expected future inflation, and some measure of real activity, and the coefficients on the different terms will depend on factors such as such as the degree of pricing power of firms, and how often firms reset their prices.

A large amount of research has focused on assessing the ability of the New Keynesian Phillips Curve (NKPC) to predict a path for inflation that is consistent with the actual data. To test whether the NKPC model is able to predict movements in actual inflation, the model-based measure of inflation is often represented as the present value of current, and expected future, costs. This representation is typically referred to as fundamental inflation. It has been shown that the fundamental inflation predicted by the NKPC tracks actual inflation fairly well using US data. One difficulty with this result, however, is that the assessment of the empirical performance of the model is often qualitative and mainly based on graphical inspection of fundamental and actual inflation. That is, the fit of the model is not evaluated statistically.

In this paper, we note that the fundamental inflation predicted by the NKPC is only a point estimate, and that its measurement is associated with uncertainties. These uncertainties arise since fundamental inflation is derived using estimates of the parameters in the NKPC, and of expectations of future costs. The object of this paper is to supplement the fundamental inflation measure with information on the uncertainties associated with its measurement. We represent these uncertainties in the form of a confidence band around the measure of fundamental inflation. This gives us an an upper and a lower limit for fundamental inflation predicted by the NKPC. By inspecting whether actual inflation falls within the bands predicted by the NKPC we can assess whether, in a given period, it is able to account for the movements in actual inflation.

On the empirical side, we present an application of our method to UK and US data. We confirm that the fundamental inflation predicted by the NKPC tracks actual inflation reasonably well for both countries. The UK measure of fundamental inflation uncertainty implies quite a narrow band and suggests that there are a number of periods where the model is not capable of accounting for movements in actual inflation. By contrast, we find that for the United States, fundamental inflation is more uncertain, casting some doubt on the empirical success of the NKPC.

Returns to equity, investment and *Q*: evidence from the United Kingdom

Summary of Working Paper no. 310 Simon Price and Christoph Schleicher

Tobin's *Q* is the ratio of the market value of a firm to the value of the firm's capital stock. The simple idea that makes it so attractive is that the larger this ratio the cheaper it is for the firm to increase the capital stock by issuing more equity. So one might expect that investment would be positively related to it, and this can be given a rigorous theoretical explanation. But it is commonly believed that, contrary to this neoclassical theory, Tobin's *Q* is of little practical use in explaining aggregate business investment. By contrast, recent evidence suggests that the user cost of capital (effectively, the equivalent to the cost of renting capital) has a statistically significant impact on investment. This is odd, because the theory in both cases is based on the same conditions, those required by firms seeking to maximise their value to shareholders. We do not attempt to resolve this empirical puzzle, but take a different approach to using the information in the data.

The value of a firm can be thought of as the discounted sum of future profits; the present-value. *Q* is therefore the ratio of this present-value to the cost of replacing capital. Standard finance theory predicts that because this present-value condition comes from future profits, *Q* should contain information about market participants' expectations of future events. The intuitive explanation is that if *Q* rises above its long-run average value, this should be an indication that either (i) future investment opportunities are expected to be good or (ii) that future investment is discounted at a lower-than-normal rate (or both). Some recent work on US data suggest that the same present-value condition relates *Q* to expected values of several financial variables such as bond yields, the ratio of debt to capital, growth in debt, and stock returns. In this paper we contribute to this debate by employing data using Bank of England estimates of the capital stock of the UK business sector.

The approach implied by standard investment theory strictly requires us to work with a marginal measure (the discounted profits relative to the cost of an extra unit of capital). Unfortunately, this can be proxied by the average (which is much more easy to measure) only under stringent restrictions which are unlikely to hold in practice, and this might explain the lack of success in some previous empirical applications. But the present-value approach employed in this paper relies on a small number of assumptions, and requires only an average value of *Q*. The main condition for the present-value framework to be valid is that average *Q* is stationary (meaning that the mean and variance of the variable in question do not tend to change over time). It is quite reasonable, theoretically, to expect this to hold. Using a battery of statistical tests for stationarity, we find evidence that this is the case. Having established this, we then look at the short and long-run relationships between *Q* and the variables it might predict, as implied by a close examination of the present-value condition. This is done in two ways.

First, we look at a system of equations to see if past values of *Q* have any additional predictive power when other lagged variables are also used to explain the data. Our results indicate that *Q* does predict the debt to capital ratio, growth in debt and investment. However, contrary to some US results, we do not find evidence that it predicts short-run fluctuations in equity returns or firms' earnings.

Second, we look at the question of whether *Q* can by itself predict variables for horizons ranging from 1 to 32 quarters, a common method in empirical finance studies. There are some well-known statistical problems arising from the fact that the tests for statistical significance are biased by the 'overlapping' nature of the data, which (among other things) causes forecast errors to be very strongly correlated between observations. We use some standard test corrections to take care of this, but we also consider some less commonly used corrections. These included 'bootstrapped' standard errors (where the uncertainty about our estimates is estimated by taking repeated samples of the original data), and a newly developed theoretical correction (derived under the standard assumption of a 'long' sample length). These different methods provide a coherent picture, in the sense that *Q* is able to predict equity returns as well as the investment to capital ratio and changes in the capital stock. In particular, as predicted by theory, *Q* is negatively related to returns and positively related to investment and capital growth, over medium and long horizons.

We therefore conclude that, at least for UK data, the common perception that *Q* is interesting from a theoretical perspective, but of little empirical relevance, is not true. In contrast, it appears to be a rich source of information about real and financial quantities.

The yen real exchange rate may be stationary after all: evidence from non-linear unit root tests

Summary of Working Paper no. 311 Georgios Chortareas and George Kapetanios

Purchasing power parity (PPP) is the hypothesis that goods will trade at roughly the same price in different countries, once adjustments have been made for exchange rates. It is usually thought of as a long-run proposition. One way of examining this is to see if the real exchange rate (the exchange rate adjusted for relative prices in different countries) tends to return to a long-run average. This is known as mean reversion, and is one of the characteristics of a 'stationary' process. The empirical literature that tests for PPP by focusing on the stationarity of real exchange rates has so far provided, at best, mixed results. The behaviour of the yen real exchange rate, of all major currencies, has most stubbornly challenged the PPP hypothesis and deepened this puzzle. The yen real exchange rate in the post-WWII era has been characterised by a trend-like appreciation. Earlier attempts to reconcile the movement of the real yen with PPP theory included consideration of behavioural breaks, but the results were disappointing. As a consequence, Japan is often considered as the typical example of PPP failure.

In this paper we provide new evidence on the stationarity of bilateral yen real exchange rates and the validity of PPP by considering non-linear behaviour; that is, the possibility that the yen real exchange rate behaves differently at low and high levels. To do so we employ a non-linear version of the widely used Augmented Dickey-Fuller test, which tests for stationarity. This extension increases the ability of the test to detect stationarity when the underlying process is non-linear. The econometric model can accommodate the possibility that an implicit 'corridor regime' exists; within this corridor real exchange rates do not converge to their average values, but once they cross the thresholds of this regime they do begin to do so. This type of behaviour is consistent with the recent theoretical models where the non-linear behaviour of the real exchange rate implies a 'band of inaction'. Our results suggest that the bilateral yen real exchange rates against the other G7 and Asian currencies were mean reverting during the post-Bretton Woods era. In particular, the bilateral yen real exchange rate against the other G7 currencies appears to be stationary over our full sample (beginning in 1960), and this result does not change when we restrict our attention to the post-Bretton Woods era (with the exception of the yen/DM real exchange rate). Thus, the behaviour of the real yen may not be so different after all, but is simply perceived to be so due to the complicated nature of its behaviour. In addition to providing support for the PPP hypothesis, our results could motivate further research aiming to explain the underlying sources of yen's non-linear behaviour.

Exchange rate pass-through into UK import prices

Summary of Working Paper no. 312 Haroon Mumtaz, Özlem Oomen and Jian Wang

Exchange rate pass-through (ERPT) is the percentage change in local currency import prices following a 1% change in the exchange rate between importing and exporting countries. A one-to-one response of import prices to exchange rate changes is known as *complete* ERPT while a less than one-to-one response is known as *partial* or *incomplete* ERPT. The rate of ERPT has important implications for the effect of monetary policy on domestic inflation as well as for the transmission of macroeconomic shocks and the volatility of the real exchange rate. As such, the relationship between exchange rates and goods prices has been studied extensively in previous work. In this paper, we focus on the pass-through of exchange rates into UK import prices, where these prices are measured for 57 industries. To the best of our knowledge, no research has been done to measure ERPT into UK import prices at this level of disaggregation, and this paper aims to fill this gap.

We use quarterly data from 1984 Q1 to 2004 Q1. Consistent with earlier studies, we find evidence for significant variation among the estimated industry-specific pass-through rates. This cross-sectional variation of pass-through rates poses an interesting problem for inference on the rate of ERPT at the economy-wide level. Our results show that ignoring this variation and simply using an aggregate import price index to estimate economy-wide pass-through rate can lead to a substantial upward bias in its measurement. Consequently, the aggregate ERPT rate can appear to be significantly higher than its true value. Using an estimation method that accounts for cross-sectional variation, we find evidence for short-run and long-run partial pass-through into import prices for the two import categories we construct using our industry-level data, namely food and manufacturing. Similarly, the economy-wide ERPT is also found to be partial, possibly reflecting the relatively large weight of manufacturing goods in UK imports. Further, we investigate the source of the cross-sectional variation in ERPT rates. Previous work on ERPT suggests that the variation of pass-through rates across industries relates to industry-specific factors such as the degree of competition, product differentiation, demand elasticities, trade barriers, inflation rates etc. For our sample, we find the industry-specific average inflation rates to be significant in explaining this variation. The final part of the paper examines whether the pass-through rates have varied across time. We find that there has been a significant decrease in the ERPT both at the economy-wide and the industry level. Our estimates suggest that this decrease can largely be explained by the increased stability of the UK economy over the past decade.

Bank capital channels in the monetary transmission mechanism

Summary of Working Paper no. 313 Bojan Markovic

Theory and empirical evidence suggest that the health and the behaviour of the banking sector can alter the way monetary policy affects inflation and output. Furthermore, a number of theoretical studies have suggested a potential role for bank capital regulation in determining bank lending decisions. Put simply, the transmission of monetary policy tightening through the banking sector is likely to be stronger when the level of bank capital approaches the minimum required by the regulator. This study assesses this 'bank balance sheet channel' using a theoretical model that extends the well-known Bernanke, Gertler and Gilchrist model of the corporate sector balance sheet channel.

The results suggest that monetary policy decisions can have a stronger effect in times when the health of the banking sector deteriorates. Banks may find it more costly to raise the fresh capital required to fulfil regulatory requirements. Moreover, the cost of raising fresh capital may rise further in economies where banks are not rated by external rating agencies, or where they disclose less information to the public, since in such circumstances potential bank shareholders may find it more costly to check the health of a particular bank. This higher cost of bank capital is further transferred to a higher cost of companies' external borrowing through an increase in loan interest rates.

This study further suggests that the impact of monetary policy can be asymmetric. An increase in interest rates is likely to lead to a fall in the value of bank capital, thus increasing the likelihood of hitting the binding capital constraint. If the latter occurs, banks have either to raise fresh capital or to reduce their loan supply. In contrast, a fall in interest rates does not produce similar effects where the additional capital is in excess of regulatory requirement.

The importance of 'the bank balance sheet channel', modelled here, can therefore vary over time. First, the bank capital constraint is more likely to bind in times of contractions (ie rises in interest rates). Here a greater need for banks to raise fresh capital coincides with an increase in the cost of it. Furthermore, the importance of the channel increases at times when the health of both the banking and corporate sectors jointly deteriorates. In contractions, companies' internal funds may dry up, and they have to rely more on external borrowing. The higher loan demand could lead to a binding bank capital constraint, which is exacerbated by the lower value of bank capital. Finally, the relative importance of 'the bank balance sheet channel' is likely to increase in periods of occasional, but large, direct shocks to banks' balance sheets. Such shocks may occur as a result of regulatory changes or structural reforms of the banking sector. Changes in the framework of bank capital regulation or an economy-wide write-off of non-performing loans are examples of such a situation.

There are a number of potential avenues for further work. Contemporary discussions about the new Basel proposals for international bank capital regulation and their potential impact on the effectiveness of monetary policy could be addressed in this framework. The analysis in this study does not however deal explicitly with the case of 'credit rationing', when banks limit their credit supply below the level of credit demand, given the same loan interest rate. In such a case the contractionary effect may be even stronger than the one proposed in this study.

Consumer credit conditions in the United Kingdom

Summary of Working Paper no. 314 Emilio Fernandez-Corugedo and John Muellbauer

It is widely perceived that credit supply conditions for UK consumers have been liberalised since the late 1970s, with implications for the housing market and consumer spending. Consumption and the housing market (with changes in credit availability likely to have contributed) were important factors in the economic boom of the late 1980s and the subsequent recession of the early 1990s.

The need for a credit conditions index (CCI), which measures credit availability other than through the level of interest rates, has been recognised in previous work on consumption. Proxies such as unsecured credit to income ratios and interest rate spreads have been used in empirical work. However, such proxies are unsatisfactory because they depend on the economic environment. This paper constructs a CCI, that, as far as possible, is free of this criticism because it controls for the effects of the economic environment. The paper constructs a CCI for households between 1976 and 2001. The index is constructed by assuming that there is an unobserved common influence (credit conditions) in each of ten credit indicators. Because this is assumed to be the same in each indicator, it is possible to back out an estimate. The history of institutional changes in the credit markets is used to guide the estimation of the CCI. Two of the ten credit indicators are aggregate unsecured debt and mortgages (secured debt). The remaining eight consist of the fractions of high loan to income and high loan to value mortgages for UK first-time house buyers split by age and regions. We argue that mortgage defaults largely arise from the coincidence of having a poor debt/equity position and experiencing cash-flow problems. So mortgage lenders limit initial loan to value and loan to income ratios to control the risk of default. We use these arguments to model the fractions of first-time buyers with high loan to value ratios and high loan to income ratios. We build on previous literature to derive specifications for aggregate unsecured and mortgage debt, although the attention to expectations and risk distinguish these models from previous work.

To ensure that, as far as possible, the CCI is not affected by the economic environment, we test and include a large set of economic controls. We start from a very general specification, so we carefully consider what theory tells us the effects of the controls should be. As far as possible, the CCI should measure credit availability, ie the supply of credit available to a typical household, once economic and demographic influences have been removed. The econometric results produce two credit condition indices. In one case, the CCI has only a direct impact on the level of credit. In the other, it works in combination with other variables, so that, for example, the influence of the real interest rate and housing wealth on debt shifts with CCI. Both indices increase in the 1980s, peaking towards the end of the decade. They fall partway back in the early 1990s, before increasing again towards the end of the sample. All equations include a common risk factor that depends upon a measure of inflation volatility, the change in the unemployment rate and a measure that is designed to capture the possibility of housing returns declining, all in the previous two years, and the mortgage possessions rate in the previous three years. At the same time, we also estimate new models for unsecured debt and mortgage debt.