

---

# Import prices and exchange rate pass-through: theory and evidence from the United Kingdom

*Working Paper no. 182*

Valerie Herzberg, George Kapetanios and Simon Price

The appreciation of sterling that began in 1996 appeared to feed through into import prices slowly, although there has been considerable downward pressure. Importers appear to have taken the opportunity to raise their margins. This is the topic examined in this paper. We begin by discussing competing theories to explain this phenomenon. We then estimate a model capturing some aspects of the process, focusing on the role of competitors' prices.

The presence of nominal rigidities has often provided a useful explanation for short-run variations in the real exchange rate. But in order to explain the persistent failure of import prices to fall fully in line with the exchange rate, we need more than simple nominal rigidity. Such persistent changes in the real exchange rate suggest the presence of pervasive market segmentation across countries. This is defined as the ability of firms to charge different prices for an identical good in different markets, or alternatively as the ability of firms to price to market (PTM). International market segmentation and imperfect competition then imply that there may be relatively little pass-through of exchange rates to import prices.

On the time series properties of the data, real exchange rate movements are well known to be volatile, but are also highly persistent. Furthermore, there is evidence that most of the real exchange rate variability is due to *traded* goods prices. Moreover, the new open economy macro models show that when such structures are embedded in complete general equilibrium models, there are profound implications for the monetary transmission mechanism. It seems that variable margins, sticky prices and (implicitly) transport costs are crucial elements in explaining the persistence in relative price changes, although not everything can be explained.

One possible implication for firms' import (and domestic) price setting is that competitors' prices affect the mark-up over marginal costs, and it is on this that we focus in this paper. Some of the factors supporting pricing to market may also introduce non-linear responses to exchange rate shocks. That is, small changes in the exchange rate may leave the prices of imported goods unchanged in sterling terms, but large changes that cross a 'threshold' may trigger an adjustment. However, although this kind of behaviour is

plausible for firms, the case for aggregate effects is less certain. Aggregation may lead to smooth non-linearities, however, and we discuss ways of testing for this.

We establish that there is evidence for pricing to market and a role for competitors' prices by estimating a mark-up equation with a role for UK prices. The mark-up is over the major six (M6) countries unit labour costs, which implies that the underlying technology is Cobb-Douglas. The mark-up is affected by domestic demand, measured by a measure of capacity utilisation, and by domestic prices. The Johansen results suggest that there is no cointegration, but the method is known to be sensitive to the maintained assumptions and specification. Given weak exogeneity, single-equation methods are a robust method, and we are able to show that the explanatory variables are indeed weakly exogenous to the long-run relationship. The equation gives a weight of 0.36 to labour costs and 0.64 to UK prices, so PTM appears to be dominant. It is possible that some import pricing may be characterised by purchasing power parity (PPP), although the adjustment coefficients in the system's individual equations suggest this is not a major factor, so we think of our estimates as offering an upper bound on the degree of PTM.

There are potential identification problems. In particular, if PPP held then our equation might conflate the firm's mark-up equation with the PPP relationship. But there is no evidence for multiple long-run relationships. And the evidence from the adjustment coefficients supports PTM as well, as the direction of causality indicated by the adjustment coefficients is from the long-run relationship to import prices. Thus we are confident that there is indeed some pricing to market in UK imports.

By contrast, there is no evidence for non-linearity, either in extreme threshold behaviour or smooth-transition models. We base this conclusion on the results of general tests that are powerful against a range of non-linear alternatives: a specific on-off threshold model where there is an abrupt transition between regimes; a simple spline model allowing differential adjustment speed at positive and negative disequilibria; and a flexible smooth-transition model that allows for continuous variation between regimes, but still nests the extreme case.

---

# Capital flows to emerging markets

*Working Paper no. 183*

Adrian Penalver

Capital flows to emerging market economies have historically occurred in cycles of enthusiasm and despair. During the upswing, confidence is high and countries may overborrow relative to the set of profitable investment opportunities, thereby creating the conditions for a financial crisis and capital outflow. Countries might be better off if they borrowed at a steadier rate and avoided these cycles in capital flows. If borrowing exceeded this optimal rate, policy-makers could take steps to restrain capital inflows or promote them if borrowing fell below this rate. But what is the optimal rate of capital flows to emerging markets? Economic theory has had very little to say on the matter. To help answer the question, this paper investigates an open-economy growth model adjusted to make it appropriate for analysis of emerging market economies. This model is then calibrated using the results of a simple econometric equation and some assumptions about the other parameters. From this, estimates of optimal capital flows to a selection of emerging market economies are reported.

Two sorts of capital are used to produce output in the theoretical model. Some capital, such as factories, ships or pipelines, can be used as collateral for international loans. This is because the assets can be owned by foreign investors so that in the event the borrower defaults, an international lender can claim the collateral and recover the money. Other capital, such as human capital, cannot be used as collateral because the asset cannot be bought or sold. For example, a creditor cannot seize the education or health of a bankrupt debtor and sell it to someone else. The first sort of capital can be used to borrow money internationally, the second sort cannot.

A capital-scarce emerging market country will borrow to invest in the first sort of capital as much as it can. However, it needs to generate resources internally to invest in the second sort of capital. But citizens of the country will also want to consume now, so the growth rate is determined by the trade-off between the desire to

consume now and investing to consume more in the future. Both forms of capital are assumed to be complementary in production, so the accumulation of capital that can be used as collateral will depend on the rate of investment in capital that cannot be used for this purpose. Therefore, the rate of international borrowing can be estimated by deriving the rate of growth in capital that cannot be used as collateral.

One feature of emerging markets is that a significant proportion of the labour force does not use internationally collateralisable capital, for example those engaged in agriculture or rural industry. This paper extends a model by Barro, Mankiw and Sala-i-Martin by adding a 'traditional' sector which does not use collateralisable capital in production. Other things being equal, the larger the traditional sector, the slower the economy will grow. However, there are other fundamental factors which also determine the growth rates of emerging market economies. To help calibrate the model, an econometric equation is presented that estimates the effect of these factors. By combining the theoretical model, the econometric equation and some additional assumptions, estimates of capital flows to a selection of emerging market economies are calculated. These estimates provide a benchmark against which to compare observed capital flows. The capital flows derived from this exercise are lower than those observed over the estimation period (1988–97), suggesting that actual capital flows might have been too high. However, the results are sensitive to the parameters chosen. Therefore, larger flows than the benchmark are not necessarily a signal of overlending. They do suggest, however, that policy-makers should take a closer look at the fundamentals of the economies concerned. Substantially higher flows can be consistent with the theory, but require confidence in underlying parameter values outside the normal range. These results cannot replace judgment on the strengths and weaknesses of an economy's fundamentals, but they can suggest where these judgments need to be made.

---

# The effect of payments standstills on yields and the maturity structure of international debt

*Working Paper no. 184*

Benjamin Martin and Adrian Penalver

Financial crises appeared to become more prevalent and more severe over the 1990s. In response, policy-makers have sought mechanisms to reduce the probability of crises occurring and to limit the costs when they do occur. One such mechanism is the temporary suspension of debt payments: a standstill.

Standstills offer potential benefits in both liquidity and solvency crises. In a liquidity crisis, a standstill would play a role similar to a bank holiday in the domestic bank run case. As such, standstills could forestall a liquidity crisis, thus preventing a liquidity crisis from degenerating into a solvency crisis. In both liquidity and solvency crises, standstills pre-empt the creditor co-ordination problem by temporarily imposing a collective solution. A common criticism of standstills, though, is that they will lead creditors to lend over shorter maturities to be well placed for a 'rush to the exits' if there is a risk of a standstill being called. This would raise the proportion of short-term debt and so could increase vulnerability to a liquidity crisis rather than reduce it. It would also potentially, then, increase the cost of capital for emerging markets.

This paper develops a simple model to analyse the effects of standstills, using comparative statics between a regime with and without standstills. The three-period model comprises an emerging market debtor and risk-neutral international creditors. The debtor needs to borrow to finance production and can issue either short or long-term debt. The key assumptions of the model are that the probability of crisis increases in the level of short-term debt, that crises have costs that spill over into the next period, and that orderly crisis resolution through the use of standstills will reduce the cost of crisis. There is full information and a competitive market for funds. A standstill is depicted as an orderly rollover of short-term debt from the first period into the second period in the event of a crisis. Investors are impatient and so dislike being caught within a standstill, which is reflected in the interest rates they charge. The debtor can choose strategically to default, but this will reduce output in the following period because investors can distinguish between incapacity and unwillingness to pay. The debtor will maximise expected net output, by choosing the optimal level of short-term debt, from which the other variables are determined.

A numerical example is considered to demonstrate the intuition of the model. The level of short-term interest rates for a given level of lending is higher under standstills, reflecting investors' impatience if caught within a standstill. But long-term interest rates are initially lower under standstills, because the lower cost of crisis reduces the risk of investing in bonds. Faced with higher short-term interest rates and lower long-term interest rates, a debtor country will lengthen the maturity of its debt, which reduces the probability of crisis. This comes at a cost of lower output.

One of the main assumptions underlying the analysis in this paper is that standstills mitigate some crisis costs. Although the reason is not modelled here, this reflects a view that disorderly resolution of financial crises imposes costs on the economy through channels such as loss of market access, reputational costs, a credit crunch, disruptions to the payments system and so on. If the crisis resolution effect is strong, standstills could raise expected output compared with the no-standstills regime. If the crisis resolution effect is weak, standstills cannot improve on the no-standstills regime, because the debtor is fully disciplined in taking risks through market prices. Expected output, however, may not be the appropriate welfare measure if crises have wider social costs than forgone output. If the national authorities were prepared to trade off expected net output and the probability of crisis, then standstills could still improve social welfare.

In summary, the model looks at the implications of standstills for yields and the maturity structure of international debt. The model suggests that creditors will not 'rush for the exits' by lending over shorter maturities. Creditors will charge interest rates that reflect the risks they face. As a result, debtor countries will tend to issue longer maturity debt if they face a tilting of the yield curve. Standstills have the benefit of reducing the proportion of short-term loans and so the probability of crisis will fall. But the cost generally is lower expected output. A country considering introducing a standstills regime would have to weigh up the welfare benefits against the potential output cost.

---

# What does economic theory tell us about labour market tightness?

*Working Paper no. 185*

Andrew Brigden and Jonathan Thomas

The aim of this paper is to offer a coherent framework for examining the underlying drivers of labour market tightness, and the relationship between labour market tightness and inflation. Our motivation stems from the fact that although the phrase ‘labour market tightness’ is frequently used in the economics literature, it is rarely defined. Nonetheless, a variety of empirical evidence on labour market quantities and prices, such as unemployment and average earnings growth, is often cited as evidence of changes in the tightness of the labour market. Without a clear definition of the phrase it is difficult to evaluate the usefulness of any evidence offered; and a proper understanding of the relationship between tightness and inflation is also problematic.

In our view ‘labour market tightness’ can be defined in terms of its implications for the labour share of income. This follows from the notion that the labour market is tight (loose) when there is an imbalance between labour demand and labour supply, which will exert upward (downward) pressure on real unit labour costs, or equivalently on the labour share. Because the words tight and loose imply a degree of imbalance, we assert that the labour market can only be considered tight or loose out of steady state. This has two important implications. First, no shock can cause the labour market to become tight or loose unless it pushes the labour market away from its steady state. In practice, this is not too restrictive, since the kinds of rigidity that are present in most popular macro-models are sufficient to do this. Second, any shock that alters the steady-state value of the labour share cannot be said to have made the labour market permanently tighter or permanently looser. This is because movements in the steady-state do not involve any change in the balance between the demand for, and supply of, labour.

We use our definition of labour market tightness and its associated properties to examine the tightness implications of several popular labour market models. We start with the basic competitive model, and then work through models of efficiency wages, insider power, skill mismatch and matching frictions. A key message of this exercise is that the implications of much-cited indicators of changes in labour market tightness, such as unemployment, depend critically upon both the underlying economic shock and any market rigidities. For example, in the model of insider power a positive shock to nominal money balances leads to a tightening of the labour market that is accompanied by a decline in unemployment, which subsequently rises over time back to its unchanged steady-state value. On the other hand, an adverse labour supply shock in the perfectly competitive or efficiency-wage models leads to a tightening of the labour market that is accompanied by a rise in unemployment to a higher steady-state value.

We then turn to the relationship between tightness and inflation. By our definition, a tightening of the labour market will cause the labour share of income to rise. Since labour market tightness is a real phenomenon, it will have no implications for inflation unless the economy is subject to some form of nominal rigidity. Examples of such rigidities that could plausibly underpin a link between tightness and inflation include sticky price expectations, and restrictions on the frequency with which firms can alter prices. If such frictions are present, it is possible for out of steady-state movements in the labour share to influence inflation. A key lesson from this analysis is that any attempt to infer the relationships between labour market tightness, various market indicators of it, and inflation, requires both a clear definition of tightness and depends on the specific model of the labour market.

---

# Ready, willing, and able? Measuring labour availability in the UK

*Working Paper no. 186*

Mark E Schweitzer

The unemployment rate is the standard measure of labour availability. This research investigates broader summary measures that account for the substantial role that people who are regarded as economically inactive, can play in employment growth. Inactivity includes people who are studying, sick or disabled, taking care of family members, and retired early. While most people in these groups are less likely to start working in the next three months than the recently unemployed, they often do take up work. So they can serve to augment the unemployed pool as a source of workers, particularly as the unemployment rate declines in an expansion. This paper examines the extent to which they have done so in the current expansion.

The approach taken here is to focus on people who are actually getting jobs (including all inactive categories), rather than limiting the analysis to those who are categorised as unemployed (ie searching and available for work). The models estimate the expected number of transitions to employment from the full set of working-age non-employed, allowing the data to indicate who is relevant. If the behaviour of the inactive were unpredictable, or did not vary over the business cycle, an expanded reference group would add little to our understanding of likely labour market pressure. However, both of these conditions are rejected by the data.

The paper considers several alternative models based on different sets of explanatory variables, including the reasons for the individual's non-employment and their personal characteristics, including age, sex and education. The models are all estimated on the same set of non-employed individuals, but include various levels of information about the individuals. For example, the unemployment-rate model only includes a single indicator of the individual's status: 1 for being unemployed and 0 for all other non-employed groups. The model in that case predicts a transition rate for both groups and an aggregate transition rate for the economy as a whole. By building models in this manner, the same yardstick can be used to compare distinct groups of available labour.

The model comparisons suggest three major conclusions. First, a model of availability relying only on the unemployment rate is based on a false premise that other categories of the non-employed are considerably less relevant to labour supply. Second, models that combine information on the classification of the inactive with demographic information do best in explaining labour supply. Finally, models based on the unemployment rate tend to overstate the recent falls in amount of labour available for employment.

---

# Sovereign debt workouts with the IMF as delegated monitor—a common agency approach

*Working Paper no. 187*

Prasanna Gai and Nicholas Vause

Countries experiencing financial crises usually look to official and private sector lenders for new credit to help them meet their financing needs. Often new loans are extended, but on conditions that various domestic adjustments also be made. This is an example of principals (the creditors) offering an incentive scheme (the conditional credit) to an agent (the debtor country). Since a country faces many different creditors with different goals, the inability of creditors to cooperate can place conflicting demands on a debtor, which may affect its ability to satisfy each creditor. In environments such as debt workouts, where creditor non-cooperation is significant, the conditions placed by official creditors (often through an IMF programme) influence the lending behaviour of the private sector. This paper uses principal agent theory to examine how the design of IMF conditionality influences the behaviour of private lenders and debtors and clarifies the influences on the provision of financial support for countries in trouble.

There are three key features to the model. First, official and private sector creditors are treated as separate entities with different goals that tend to pull the debtor country in different directions. For example, if private lenders are interested only in ensuring that the debtor meets its short-term obligations, they might press for actions that raise finance quickly, such as the sale of state assets. At the same time, if the official sector ('the IMF') is promoting long-term debt sustainability, it might encourage the debtor country to pursue various economic stabilisation policies as well as structural reforms. Second, the official sector is assumed to extend

credit before the private sector. It is often the case that the debtor countries in trouble approach the IMF to assist with rollovers of credit lines. And third, the official sector observes performance measures, which are, to some degree, aligned with the actual outcomes that result from a debtor's adjustment effort. The provision of IMF credit is linked to a country's performance criteria. But since performance measures (such as ceilings on net domestic assets) are themselves surrounded by uncertainty, they cannot be too narrowly defined. Appropriate conditionality must therefore balance the controllability of a performance measure with its alignment to actual outcomes.

Since creditors are unable to observe a debtor's actions perfectly, a debtor has an incentive to side step the conditions stipulated by creditors during crisis management. We show how this '*ex post* moral hazard' is exacerbated by the lack of creditor co-operation. Our results suggest that IMF intervention in the debt workout, where it has seniority rights over its loans with respect to other types of creditors, can mitigate some of the inefficiencies due to creditor non-cooperation. Delegating the task of policy-conditional lending to an agency like the IMF leads to a more efficient outcome characterised by increased lending and rollovers by the private sector. But the ability of the official sector to do this depends critically on the focus of the IMF programme. Our findings highlight the importance of 'result-based' conditionality and mechanisms that enhance the ability of the official sector to monitor and enforce good policy behaviour and exercise leadership during debt workouts.

---

# The role of asset prices in transmitting monetary and other shocks

*Working Paper no. 188*

Stephen P Millard and Simon J Wells

This paper aims to improve our understanding of both the information contained in asset prices, and their response to monetary policy. In particular, we concentrate on the fundamental determinants of asset price movements, and rule out asset price bubbles or speculation. To be more specific, we are interested in the following questions about asset prices:

- (a) What is the response of asset prices to monetary policy?
- (b) What information do asset prices contain about fundamental shocks affecting the economy?

Although (a) and (b) are essentially different issues, our approach is to build a common framework in which to answer them. We first estimate an empirical model to characterise the data. Then, we construct a dynamic stochastic general equilibrium (DSGE) model that enables us to address our questions at a fundamental level. Using this dual approach, we are able to uncover interactions between asset prices and monetary policy that are theoretically and empirically congruent. The DSGE model is then used to uncover the interactions between key economic variables and assets in the face of fundamental economic shocks that are difficult to identify empirically. This allows us to draw robust conclusions about the role of and informational content of fundamental asset prices and their role in the transmission mechanism.

The empirical model is a fairly standard vector autoregression (VAR) that identifies a monetary policy shock. The empirical model serves two purposes: by plotting the response of asset prices to monetary policy shocks it gives an entirely empirical answer to question (a), and it provides a benchmark with which to evaluate a theoretical model. We find that output falls after a contractionary monetary shock and exhibits a 'hump-shaped' response. Base money and the price level also fall. But, as is often found in this type of analysis, the effect on the price level is small. There is a short-lived rise in the exchange rate, which then follows a UIP path and gradually depreciates back to base. Short-run nominal interest rates rise after the shock. However, we find a small but significant fall in the long rate. This implies a clockwise rotation in the yield curve. House prices fall, but more quickly and by a larger amount than the general price level. There is also a short-lived fall in equity prices.

The baseline for the theoretical model is an open-economy 'Consumption CAPM' model, based on a representative rational

agent who can hold a portfolio of various assets: domestic real and nominal bonds, foreign nominal bonds and shares in domestic firms. Agents select their portfolios to maximise the present value of their lifetime utility. The model is fairly standard with consumers exhibiting habit persistence over their consumption of traded and non-traded goods and housing services. In order to obtain housing services, they combine time spent in household production with the existing housing stock. Their derived demand for housing together with exogenous housing supply generates interesting dynamics for house prices. Our interest in house prices is motivated, in part, by empirical studies that have shown house prices to be a useful indicator for inflation. Firms combine labour and capital services to produce both traded and non-traded goods. We assume that investment decisions in the non-traded sector are subject to convex 'costs of adjustment'. The model is calibrated for the United Kingdom. By shocking variables in the model, and plotting the response of asset prices to shocks, we can answer our two questions at a fundamental level.

We first examine the responses of variables to an exogenous monetary policy shock in the theoretical model. We find that the theoretical model is able to produce responses qualitatively similar to those uncovered from the data, although we also find some differences. We take the general congruence between model and data to be an encouraging sign.

We then use our model to show how a given fundamental shock may imply a unique pattern of asset price movements in the periods immediately after the shock. Therefore, observing patterns of asset prices and comparing them with the movements implied by our model might reveal the nature of shocks currently hitting the economy. There are, however, several important reasons why this information should only be used tentatively and to corroborate other evidence. First, the results presented in the paper are dependent on the monetary policy response to shocks. That is, they depend crucially on the monetary policy rule that we assume being a reasonable characterisation of the monetary reaction function, and that this is fully known by market participants. Second, asset prices often move for reasons not obviously related to economic fundamentals; we should be careful not to assume that any movement in asset prices is driven by fundamentals.

---

# Modelling investment when relative prices are trending: theory and evidence for the United Kingdom

*Working Paper no. 189*

Hasan Bakhshi, Nicholas Oulton and Jamie Thompson

Recent research in the United States has shown that aggregate economic models fail to explain the investment boom in real plant and machinery in the second half of the 1990s. In contrast, a disaggregated modelling approach does much better. This appears to reflect two factors. First, aggregate models do not capture the increase in replacement investment associated with compositional shifts in the capital stock towards shorter-lived assets, such as computers. Second, aggregate models invariably find little or no role for the real user cost of capital, so they understate the positive effects of falls in the relative price of computers on investment in computers.

The United Kingdom also experienced a boom in real plant and machinery investment in the second half of the 1990s. But undertaking similar research is beset with difficulties in the United Kingdom, not least because of the relative paucity of disaggregate investment data in the published National Accounts. In this paper, we carefully construct a data set for the United Kingdom that is consistent with the National Accounts. We then use these data to investigate the ability of different investment equations to account for the UK investment boom in plant and machinery. We report results similar to those for the United States. In particular, the traditional aggregate modelling approach completely fails to explain the investment boom in plant and machinery in the second half of the 1990s.

Our analysis consists of two main elements: a theoretical section setting out the relationship between aggregate and disaggregated approaches to modelling investment; and an empirical analysis setting out our econometric results.

In our theoretical analysis, we first derive the relationship between firms' desired capital stocks and the real user cost of capital, as predicted by standard economic theory. We show how that relationship breaks down in the presence of a trend decline in the relative

price of investment goods. Such a trend has been a particularly important feature of investment in recent years. In contrast, we show that well-specified relationships exist at the disaggregated level.

Our empirical exercise involves using time series cointegration methods to model investment at disaggregated and aggregate levels. We compare the ability of the two approaches to explain the boom in plant and machinery investment. Recognising that cointegration techniques can have low power, particularly in small samples, we further evaluate the comparative performance of the two approaches by conducting out-of-sample forecasting exercises.

In all cases, our empirical results support the theoretically superior disaggregated modelling approach. First, compositional shifts in the capital stock towards shorter-lived computer assets appear to have been important in the United Kingdom too in the second half of the 1990s. That explains some, though not all, of the inability of the aggregate model to explain the investment boom. The second factor behind the strong investment growth has been a decline in the relative price of computers. Echoing findings for the United States, we find that firms' investment in computers appears to be highly sensitive to falls in the real user cost of computers. And interestingly, our models suggest that the increase in the size of firms' computer capital stocks in the second half of the 1990s is fully accounted for by the sharp falls in the real user cost of computers.

Given the great uncertainties surrounding measures of the real user cost of capital and the price of investment goods in particular, we investigate the sensitivity of our results to alternative measures of the real user cost of capital. The paper's results are reassuringly robust. Overall, they provide strong support to attempts to model and forecast investment at the disaggregated level.

---

# What caused the 2000/01 slowdown? Results from a VAR analysis of G7 GDP components

*Working Paper no. 190*

Vincent Labhard

The recent slowdown in the world economy has rekindled interest in the major shocks affecting the business cycle, both at a global and country level. While this interest is not unusual at the current phase in the cycle, there are other factors that explain why this time it has been more intense than usual. First, the recent slowdown came at the end of one of the longest expansions on record. Second, it seemed to have affected several countries at the same time, so one question is whether the slowdown was due to common shocks. And third, there was a perception that the international transmission of shocks has changed.

Using a method first developed by Olivier Blanchard, this paper conducts an analysis of the shocks to GDP components, which in our case include private and public consumption, residential, business and government investment, exports, imports and changes in inventories. Such an analysis has two main benefits. First, it provides a counterpart to stories cast in terms of component developments, for example the role of revisions to expected future profitability, especially in the information, communications and technology sector, or the role of the millennium changeover. Second, it provides a more detailed picture of developments during the slowdown than can be obtained by looking only at GDP, and thus a useful cross-check on work aimed at identifying the shocks affecting GDP.

The analysis is based on a vector autoregression of GDP components, which is used to account for their interrelationships and to provide series of corresponding innovations. There are potentially other variables explaining the behaviour of GDP components, but using only components data has the advantage of capturing empirical regularities in a parsimonious set-up. The innovations are then used to extract component-specific shocks, which form the basis for the analysis. These shocks capture the movements in the components that are not explained by the components' history and exclude the factor common to all components. So the shocks are a catch-all of a range of potential structural factors. This analysis is applied to the G7 countries individually and as a group, thereby extending previous research to a cross-sectional dimension. The paper obtains estimates of the shocks during the slowdown in 2000/01, the expansion preceding it, and the previous slowdown in 1990. A second set of estimates of the shocks during 2000/01

explicitly takes into account the stance of monetary policy and the oil price.

The estimates indicate that there were shocks to several components and several countries during 2000/01. While some of the shocks were similar across the three largest G7 economies, consistent with the perception of a highly synchronised slowdown, other shocks were more country-specific. For example, there were differences in the shocks affecting Germany relative to the other countries of the euro area. Among the components, the largest and most persistent shocks in 2000/01 affected business investment, inventories and net trade. There were also large shocks to private sector consumption, but these occurred mainly in the early and late stages of the slowdown. The pattern of shocks during the preceding expansion was much more subdued. Though consumption shocks did play a role, for example, they were smaller in size and showed less persistence. More generally, the paper finds that shocks were less persistent and on average smaller (due to smaller size and offsetting signs) during the expansion, especially over a longer period. While this need not be significant, it is consistent with the common perception that expansions tend to be longer and have a slower pace than contractions.

There are several differences between the shocks during 2000/01 and those during the previous slowdown in 1990. These relate to the overall balance of shocks (which was negative for much longer in 2000/01), the major shocks (despite the shock in 2001 Q3, shocks to private sector consumption seem to have been more persistent in 1990) and the extent to which shocks were correlated across the G7 countries (while some shocks were similar for the United States, Japan and Germany in 2000/01, the United States experienced a specific pattern of shocks in 1990). Finally, when the analysis also accounts for the stance of monetary policy and oil prices in 2000/01, it appears that about half of the balance of shocks can be attributed to these factors. Their impact though varies over the period of the slowdown. Whereas they contributed considerably in the early stages of the slowdown, they became less important as the slowdown wore on, and eventually started to contribute to the recovery, a reflection of the declines in interest rates and the oil price that was under way at that time.

---

# Endogenous price stickiness, trend inflation, and the New Keynesian Phillips curve

*Working Paper no. 191*

Hasan Bakhshi, Pablo Burriel-Llombart, Hashmat Khan and Barbara Rudolf

During the 1990s the UK economy, along with several other industrialised countries, has moved to a low-inflation environment. The aim of this paper is to examine the implications of this structural change for firms' pricing decisions. We look in particular at the Calvo model of price-setting that underlies optimising small structural models and the New Keynesian Phillips curve (NKPC) in the presence of positive trend inflation. The Calvo model assumes that firms' prices are sticky, but the timing and frequency with which firms adjust prices is exogenous and constant. In recent years this model has provided an important framework for examining inflation dynamics, monetary policy rules and stabilisation policies.

In this framework it is typically assumed that trend inflation (or steady-state inflation) is zero. This assumption is particularly restrictive when examining the effect of changes in the inflation environment for firms' pricing decisions. For example, in a moderate to high trend inflation environment, firms with fixed nominal prices experience larger erosion in their relative prices and are likely to reset their prices more frequently. Hence, inflation dynamics should differ from those in a low-inflation environment. In the Calvo model, however, firms do not choose when to change prices. As trend inflation rises this assumption becomes increasingly unrealistic. Consequently, it is unclear what the upper bound of trend inflation is, below which the Calvo price-setting assumption is a good approximation, and how the structure of the NKPC—that has come under serious empirical scrutiny as a model of inflation dynamics—is affected with positive trend inflation. The answers to these questions are important when applying this framework to examine the implications of structural changes for inflation dynamics in the United Kingdom.

We consider an optimising model with Calvo price-setting that is most suitable for monetary policy analysis—one with both nominal and real rigidities (or equivalently, 'strategic complementarities')—and examine the consequences of positive trend inflation for the NKPC. We build on earlier work that also considers positive trend inflation, but ignores real rigidities. In the presence of real rigidities (for example, firm-specific factor inputs, input-output linkages, cyclical (desired) mark-ups, variable capital utilisation), firms are reluctant to adjust relative prices in response to demand shocks and consequently amplify the effects of nominal rigidities.

The Calvo assumption that firms cannot choose the timing and frequency of price changes places an upper bound on the trend inflation rate for which the model can be solved. For standard calibration, this upper bound is influenced by the interaction of nominal and real rigidities. The paper also

examines the implication of positive trend inflation for the slope of the NKPC. In an environment in which firms' price-setting behaviour is also influenced by the monetary policy regime that determines the trend inflation rate—ie firms adjust their nominal prices more frequently at higher trend inflation rates to avoid the erosion of their relative prices—this intuitively appealing extension of the Calvo price-setting has important bearing on both the upper bound and the slope of the NKPC.

Our results show that the Calvo price-setting model is well defined under annualised trend inflation rates of 5.5% or lower. Above that bound, the model implies that firms should stop production completely. This upper bound is below the average actual inflation rate for the United Kingdom over the period 1960–2000. For several other countries, the upper bound is also below the average actual inflation rates during the 1970s and 1980s. When strategic complementarities are ignored, as in earlier work, the upper bound is approximately 13%. The paper shows that, over the range of trend inflation for which the model is defined, the slope of the NKPC rises when trend inflation falls. This implies that a rise in demand pressure has a larger effect on inflation when the economy is in a low-inflation environment than when it is in a high-inflation environment. This feature sits oddly with the stylised facts and conventional wisdom that Phillips curves are flatter in a low-inflation environment.

These results can be explained intuitively by firms not choosing when to change their prices in the Calvo model. The presence of trend inflation makes firms more concerned about the future erosion of their mark-ups (and hence losses in profits). In other words, their effective discount factor rises towards unity as trend inflation increases (ie they care more about the future) and consequently their current mark-up is relatively less important. The constraint that discount factors cannot exceed unity places the upper bound on the trend inflation rate for which the model can be solved. Because the current mark-up is less important, the current output gap has a smaller effect on inflation in the NKPC.

The main conclusions from this research may be summarised as follows: (i) the low upper bound on trend inflation makes the exogenous Calvo price adjustment framework a less appealing description of how firms set prices, even in low-inflation environments; (ii) endogenous price stickiness, or more generally state-dependent price stickiness, that allows firms to choose the timing of their price adjustment to shocks, not only avoids the limitations of the standard model, but also provides a useful and promising extension of the standard Calvo model to examine structural changes.

---

# Capital stocks, capital services, and depreciation: an integrated framework

*Working Paper no. 192*

Nicholas Oulton and Sylaja Srinivasan

This paper presents an integrated framework to measure capital stocks, capital services, and depreciation. Much of the difficulty of deriving good measures of aggregate capital, whether stocks or services, derives from two empirical facts. First, the relative prices of different types of asset are changing. Second, the pattern of investment is shifting towards assets with shorter economic lives. So we cannot treat capital as if it were composed of a single homogeneous good. To some extent, these two facts are aspects of the shift in the pattern of investment towards ICT assets. The relative prices of these assets are falling rapidly and their economic lives are much shorter than those of most other types of plant and machinery.

The wealth concept of capital, while appropriate for some purposes, is not the right one for a production function or for a measure of capacity utilisation. For the latter purposes, we need a measure of aggregate capital *services*. The volume index of capital services (VICS) answers this need. In principle, the VICS measures the flow of capital services derived from all capital assets, of all types and all ages, in a sector or in the whole economy. The main difference between the VICS and wealth-type measures of capital is in the aggregation of different types and ages of assets. In the VICS, each item of capital is weighted by its rental price. The rental price is the (usually notional) price the user would have to pay to hire the asset. In wealth measures of the capital stock each item is weighted by the asset price. A VICS measure gives more weight than a wealth measure to assets like computers and software for which the rental price is high in relation to the asset price.

We review the theory of, and empirical evidence on, depreciation. The assumption that depreciation is geometric greatly simplifies the theory and seems consistent with the facts. We also consider whether this assumption is appropriate for assets like computers, which do not suffer much from physical wear and tear, but have very short lives due to 'obsolescence'. Though, in principle, our framework encompasses obsolescence, in practice depreciation rates may be somewhat overstated owing to failure to control fully for quality change.

We adopt the geometric assumption in our empirical work for the United Kingdom. Because of the uncertainty about asset lives and the pattern of depreciation in the United Kingdom, we calculate wealth and VICS measures under various assumptions. We test the sensitivity of our results in three ways. First, we compare results using both US and UK assumptions about asset lives. Second, we compare results based on a coarse breakdown of assets into four types only, with results derived from a breakdown that distinguishes computers and software separately. Third, we compare the effect of US versus UK price indices for computers and software. Our main findings for wealth and the VICS are as follows:

1. Using the conventional National Accounts breakdown of assets into buildings (excluding dwellings), plant and machinery, vehicles, and intangibles, the *growth rates* of wealth and the VICS are insensitive to variations in depreciation rates (ie,

asset lives). In these experiments the rates for each asset are assumed constant over time.

2. However, the *level* of wealth is quite sensitive to variations in depreciation rates.
3. Still sticking with the conventional asset breakdown, wealth and the VICS grew at similar rates over the period as a whole. In the 1990s, the gap between the two measures widened a bit, with the growth rate of the VICS higher.
4. The effect on the estimates of separating out computers and software is complex. First, larger differences appear between the growth rates of the VICS and wealth. Second, the growth rate of wealth tends to be slower, though that of the VICS is not necessarily faster. But under the assumptions closest to US methods, the growth rate of the VICS is raised, relative to the VICS with computers and software included with other asset classes.

These results suggest that the treatment and measurement of investment in computers and software is an empirically important issue. The relative price of these assets has been falling, so it is in principle correct to separate them out explicitly—and it matters in practice. The conclusions about the growth rates of both the VICS and wealth turn out also to be sensitive to the price index used for computers and to how the level of software investment is measured.

We also estimate aggregate depreciation (capital consumption) for the same range of assumptions. We study the sensitivity of the aggregate depreciation rate and of the ratio of depreciation to GDP to the assumptions, and compare our estimates with ones derived from official data.

1. Using the conventional asset breakdown and our assumptions about depreciation rates at the asset level, there is no tendency for the aggregate depreciation rate to rise over the past two decades.
2. Separating out computers and software has less effect than expected: even the use of US methods raises the aggregate rate only slightly, and again there is no sign of an upward trend. The reason is that, even by 2000, the share of computers and software in wealth was only about 4% in the United Kingdom. By contrast, the aggregate depreciation rate in the United States has trended smoothly upwards since 1980, illustrating the much greater scale of ICT investment in the United States.
3. Assumptions about asset lives have a large impact on the estimated ratio of depreciation to GDP. The UK National Accounts measure has been drifting down steadily since 1979. In 2001 it stood at 8%, but using shorter US asset lives and the conventional asset breakdown, the ratio was over 10%. Separating out ICT assets and using US methods, the ratio rises to nearly 13%, similar to that in the United States. In neither country was there any upward trend in the ratio, except perhaps in the past couple of years. The reason is that, although the quantity of high-depreciation assets has been growing faster than GDP, this has been offset by their falling relative price.