Investment adjustment costs: evidence from UK and US industries

Summary of Working Paper no. 332 Charlotta Groth and Hashmat Khan

If wages and prices were perfectly flexible, and if labour and capital could move costlessly between firms and sectors, the economy would always operate at potential. In this case, large fluctuations in output, consumption and investment would not be observed. But from the past we know that some of these variables exhibit large fluctuations over the business cycle. To understand these movements, it is important to acknowledge the presence of frictions in the economy, that prevent prices and the factors of productions — labour and capital — from adjusting in response to shocks.

The literature has recognised the importance of both nominal and real frictions. Nominal frictions arise when wages and prices are sticky and therefore do not respond to changes in the economic environment. These types of frictions have been stressed in the New Keynesian literature, and give rise to the well-known Phillips trade-off between inflation and some measure of real activity. Real frictions prevent labour and capital from costlessly adjusting in response to changes in the economy. As an example, consider a firm that wants to increase its stock of capital, to be able to meet an increase in demand. In addition to the cost for buying new equipment, it may also need to spend resources on physically installing the capital, training labour and reorganising the production process, to make full use of the capital. These types of costs prevent firms from costlessly adjusting the level of capital. In turn, this means that firms will only slowly respond to shocks that alter the optimal level of capital, since it may prove costly to adjust capital in response to short-lived changes in economic conditions.

Frictions to adjusting the level of capital are common in models of the business cycle, to better replicate and explain economic fluctuations. But there are some shortcomings with these models. For example, they fail in generating the hump-shaped response of output, investment and consumption that is typically observed after a monetary policy shock — an unexpected change in the stance of monetary policy. They are not able to account for the volatility of asset returns over the business cycle. And they are not able to match the response of wages and hours worked in response to fiscal shocks. For this reason, recent studies instead introduce a friction to changing investment, instead of capital, into models of the business cycle — a so-called investment adjustment cost. This friction prevents investment quickly responding to changes in economic conditions. By introducing this friction, the performance of business cycle models are

improved along a number of dimensions, such as those discussed above.

Investment adjustment costs therefore appear to have important implications for understanding the aggregate dynamics of the economy. It is, however, unclear whether there is empirical support for these types of costs at the firm or industry level, or whether they are largely an *ad hoc* friction, introduced to better match aggregate data. Some motivations have been made for these types of costs — they may proxy delays in investment planning, or inflexibility in changing the planned pattern of investment. While this interpretation is appealing, so far no attempt has been made to estimate investment adjustment costs directly at a disaggregated level. In comparison, a large body of literature has estimated capital adjustment costs using disaggregated data. The disaggregated approach is also extensively used to assess evidence on other important frictions in the economy.

In this paper we conduct an empirical assessment of investment adjustment costs and investigate whether industry-level data provide support for this cost structure. We use industry data for both the United States and the United Kingdom, and estimate a theoretical model for capital and investment under different assumptions of the adjustment cost structure. In particular, we consider a model which is a weighted average of the investment and the capital adjustment cost model, and obtain industry-specific estimates of the relevant parameters in the adjustment cost function. The main result is that the relative weight on the investment adjustment cost model turns out to be close to zero, for all industries, in both countries. In other words, industry data do not support the investment adjustment cost structure and instead favour the traditional capital adjustment costs.

We also estimate a constrained model which imposes the investment adjustment costs on the data. Based on the estimated parameters from this model, we are able to quantify the importance of the investment adjustment cost friction. We compare this estimate to those typically obtained in aggregate models of the economy. Our results suggest that at the industry level, the friction arising from investment adjustment costs is significantly smaller than that assumed at the aggregate level. From this, we conclude that from a disaggregated empirical perspective it remains difficult to motivate and interpret the investment friction considered in recent macroeconomic models.

Labour market institutions and aggregate fluctuations in a search and matching model

Summary of Working Paper no. 333 Francesco Zanetti

It is recognised that the labour market plays an important role in the assessment of the economy. The value of labour accounts for two thirds of the total value of goods and services produced in the economy. That makes labour costs a crucial influence on most firm's production and pricing decisions, and, therefore, on the dynamics of inflation and other important macroeconomic variables. This paper explores the influence of some key institutional features of the labour market on aggregate fluctuations in real quantities like output and unemployment, and inflation. It assesses their quantitative implications by studying the effects of unemployment benefits and firing costs. Unemployment benefits are modelled as payments that accrue to workers after separations from jobs, while firing costs are modelled as firing taxes that firms pay when a worker is dismissed. It is widely thought that the best approach to macroeconomics is to use a general equilibrium approach, where the evolution of the economy over time is fully integrated into the model, and the uncertain ('stochastic') nature of the world is explicitly recognised. These are known as dynamic stochastic general equilibrium (DSGE) models. This paper uses a DSGE model characterised by search and matching frictions in the labour market and nominal rigidities in the goods market, a relatively new approach.

Results suggest that an increase in firing costs decreases the volatility of output, unemployment, employment and flows

both into and out of employment, while the volatility of inflation, real wages and labour market tightness all increase. The presence of firing costs affects the intertemporal employment decision of firms, since an increase in current employment exposes firms to future firing costs. This induces firms to decrease lay-offs and hiring, leading to higher unemployment duration and lower unemployment incidence. Since quantities are more costly to change and disturbances affect a lower number of jobs, firms adjust to shocks through prices, changing them aggressively. Hence, inflation becomes more volatile.

An increase in unemployment benefits has the reverse effect. The volatility of output, unemployment, employment, and flows in and out of the labour market increases, while the volatility of inflation, real wages and labour market tightness decreases. Higher unemployment benefits make unemployment less painful for workers, causing the duration and flows into unemployment to increase. Since workers have an incentive to stay out of employment as long as they are eligible for unemployment benefits, and shocks displace a larger number of jobs, the volatility of labour market quantities increases. Firms find it more convenient to adjust the employment level in response to shocks, so that they are less likely to adjust their prices in response to disturbances. As a result, inflation volatility decreases.

Using copulas to construct bivariate foreign exchange distributions with an application to the sterling exchange rate index

Summary of Working Paper no. 334 Matthew Hurd, Mark Salmon and Christoph Schleicher

Option contracts give the right, but not the obligation, to buy or sell a financial asset at a predetermined price, known as the 'strike price'. As such, the value of an option depends on the likelihood that its holder will exercise this right. Therefore, option prices contain information about the probability that market participants attach to different outcomes of future asset prices. A common way to summarise this information is by estimating probability distributions of future asset prices implied by option prices.

Option contracts are traded for a wide range of currencies. Nevertheless, for the purpose of monetary policy, the effective exchange rate index (ERI) is more relevant than individual bilateral exchange rates. The main purpose of this paper is to develop a method that estimates option-implied distributions for effective exchange rates.

The sterling effective exchange rate is a weighted average of a large number of sterling bilateral exchange rates. However, it can be reasonably approximated by a function of only two exchange rates — the sterling-euro and the sterling-dollar bilaterals. The distribution of the sterling ERI can then be modelled as a function of the joint distribution of the sterling-euro and sterling-dollar exchange rates. The joint distribution describes the probability of all possible joint outcomes of the two exchange rates.

We observe option prices on the sterling-euro and the sterling-dollar exchange rates and use them to compute individual probability distributions for the two exchange rates. In order to link these two individual distributions to a joint distribution we make use of a so-called copula. Copulas are functions that link probabilities of individual events ('it will be cloudy tomorrow', 'it will rain tomorrow') to those of outcomes of joint events ('it will be cloudy and it will rain tomorrow'). In the context of this paper, they join two one-dimensional distributions to create one two-dimensional distribution. Copulas are useful, because they provide a very general description of dependence patterns.

A potential problem is the fact that there are a very large (in fact infinite) number of copulas to choose from. We overcome this problem by imposing a no-arbitrage condition between the joint distribution of the sterling-euro and sterling-dollar exchange rates on the one hand, and the univariate distribution of the euro-dollar cross-rate on the other hand. Observed option prices need to satisfy this condition, because otherwise they would present an opportunity to make a profit in excess of the risk-free rate without taking on any risk. Our no-arbitrage condition is a generalisation of the standard triangular no-arbitrage condition between any two spot exchange rates and their cross-rate. This narrows the choice of the copula function to those that are consistent with no-arbitrage.

In an empirical application we derive option-implied distributions of the sterling ERI on a daily basis between 2000 and 2005. We show that the distribution has seen considerable variation during this time. We also show how we can compute distributions of the sterling ERI that are conditional on movements in the euro-dollar exchange rate. This allows us to gauge the sensitivity of the sterling ERI to changes in the cross-rate of the United Kingdom's main trading partners.

In a second application we show that our method can be simply modified to calculate prices for options on exchange rate indices. We show that, contrary to standard models, the copula-based model generates a smile effect: options with strike prices that are further away from the current level of the effective exchange rate are relatively more expensive.

Business cycle fluctuations and excess sensitivity of private consumption

Summary of Working Paper no. 335 Gert Peersman and Lorenzo Pozzi

When consumers can freely lend and borrow on capital markets, aggregate private consumption should only react to changes in permanent income. Previous scientific work, however, finds that total consumption growth in the economy is determined by the growth rate in total disposable income. An important interpretation of this observation is that a fraction of the consumers in the economy is having a hard time obtaining credit. We say that these consumers are liquidity constrained. Therefore, when confronted with a higher income, these consumers tend to spend the additional amount instead of saving it. Another part of the consumer population does not face difficulties obtaining a loan and is therefore able to consume as much as it can. When confronted with a higher income these consumers do not necessarily consume the additional amount: they save it.

In this paper we investigate whether the impact of disposable income growth on consumption growth is higher during recessions than during expansions, ie whether during recessions there is a higher number of consumers who spend their disposable income. We find that this is the case. Our finding is based on a data set for the US economy that covers the period 1965–2000.

From a policy point of view, our findings suggest that the impact of policy changes that affect disposable income is very likely to have greater effects during recessions than during

expansions. Our study is motivated by theoretical results found in previous work where it is argued that during recessions liquidity constraints faced by consumers are more severe than in expansions. The reason is that the worsening of households' balance sheets in a recession decreases the possibility of consumers financing their expenditures through accumulated wealth. This raises the demand for credit. At the same time however the higher monitoring and contract enforcement costs faced by banks during recessions increases the cost of banks to give loans and therefore diminishes the credit supply. Our observation that consumption growth depends more heavily on disposable income growth during recessions thus supports previous theoretical results.

In our study we revisit an issue that was investigated in previous studies, namely the possibility that, over time, the fraction of liquidity-constrained consumers has decreased. In previous work it has been suggested that financial liberalisation and the development of credit markets that has occurred in the United States (especially during the 1980s) may have reduced the numbers of consumers that are liquidity constrained. We test this hypothesis by looking at whether the impact of disposable income growth on private consumption growth has fallen over the period 1965–2000. We find that it has not, suggesting that the average number of consumers that are liquidity constrained has not decreased.

A state space approach to extracting the signal from uncertain data

Summary of Working Paper no. 336 Alastair Cunningham, Jana Eklund, Christopher Jeffery, George Kapetanios and Vincent Labhard

Most macroeconomic data are uncertain — they are estimates rather than perfect measures. Measurement errors arise because data are typically based on incomplete samples. And they arise because many variables — for example, in-house software investment — are not easily observable; necessitating the use of proxies. Such uncertainty poses challenges for both forecasting and economic analysis. Where it is material, economists must decide how much weight to place on apparent 'news' in the published data. But how can the extent of the problem be judged and what can be done about it?

One symptom of data uncertainty is the propensity of statistical agencies to revise their estimates in light of new information (bigger samples) or methodological advances (better proxies). In the United Kingdom, the National Accounts are subject to a rich revisions process and as a result, the scale of the ensuing revisions may give a clear indication of the extent of data uncertainty in the past. And to the extent that past revisions give a good guide to the likely scale of revisions in the future, they can also be used to gauge the uncertainty associated with the latest data.

Recognition of this uncertainty leads naturally to a probabilistic view of the past. Estimation of a confidence interval around the official published data is a first step; giving an indication of the potential scale of revisions. Going further, economists can gather additional evidence about the current economic conjuncture; using that evidence to assess the likely impact of future revisions on the profile of growth.

Treating uncertain data in this way is neither new nor unique to the Bank. A 2004 study by the Statistics Commission concluded that 'the main users of the [official] statistics knew that revisions should be expected, understood the reasons for them, and were able to make some allowance for them when taking important decisions'. However, most attempts to allow for potential revisions are informal. Approaching the issue more formally can add rigour to the exercise of combining such diverse source of information — this sort of exercise is known as a 'signal extraction problem'.

This paper describes a formal model of uncertain (revisable) data that can be used to extract the signal from uncertain data. The model draws on the experience of past revisions to proxy the uncertainty surrounding the latest vintage of the official data published by the Office for National Statistics. It estimates how far and in which direction to update preliminary estimates using past patterns of revisions, alternative indicators (such as business surveys) and time-series properties of the data. The model's output is an estimate of the 'true' value of the variable of interest that can be used as a cross-check of the latest published data, or even to substitute for those data in any economic applications.

In using the model to predict the cumulative impact of revisions, economists should, however, be alert to a number of caveats. In particular, the model relies on past revisions being a good indicator of current uncertainty. It is, however, possible that revisions may become less predictable in the future. For example, successful delivery of the Office for National Statistics' Statistical Modernisation Programme will enable faster balancing of National Accounts data from differing sources and facilitate internal reviews of collation procedures. And some significant methodological revisions in the past such as the introduction of the ESA 95 accounting framework - may not be representative of current uncertainty. It is also quite possible that alternative indicators that have provided a good mapping to mature ONS data in the past will offer a worse indication in future - for example if the sample of respondents to a particular business survey becomes unrepresentative.