Network models and financial stability

Summary of Working Paper no. 346 Erlend Nier, Jing Yang, Tanju Yorulmazer and Amadeo Alentorn

Systemic risk is a key concern for central banks charged with safeguarding overall financial stability. Systemic risk arises when there is the potential for multiple banks to fail and to impose costs on the financial system and ultimately on the economy as a whole. The costs of systemic failure have been estimated to be large.

A number of determinants of systemic failure, including the role of government safety nets have been studied extensively in the literature. However, relatively little is known about how the structure of a banking system, including the degree to which banks are connected to each other through bilateral exposures may affect its susceptibility to systemic breakdown. How does the size and distribution of exposures between banks determine the resilience of the system as a whole? How does the potential for interbank exposures to transmit shocks from one bank to another interrelate with the aggregate amount of capital available to cushion shocks? In the presence of interbank exposures, are more concentrated banking systems with a small number of large banks, more or less susceptible to systemic breakdowns than systems that comprise a large number of smaller banks? And, are 'tiered' systems, where a small number of first-tier banks coexist with a fringe of smaller banks, more or less susceptible to systemic breakdown than systems that are more uniform?

This paper studies these questions drawing on recent advances in the study of networks. A banking system is represented by a set of nodes (banks) that are connected by directed links (interbank exposures) with a certain predefined probability. The 'weight' of these links (the size of interbank exposures) determines the capacity for losses to flow from one bank to another. Capital and deposits are introduced as the first and ultimate recipients of any losses incurred. In this set-up we simulate the extent of contagious (knock-on) defaults arising from losses being transmitted through interbank exposures for a wide variety of banking systems that differ in their underlying structural characteristics.

The analysis suggests that increasing the size of interbank liabilities tends to increase the risk of knock-on default relatively sharply once the size of liabilities moves beyond a certain threshold. Moreover, this is the case even if banks hold capital against such exposures in the same way as they hold capital against other credit exposures. This means that capital requirements alone may not be adequate to protect the system against knock-on defaults. Further, the simulations show that the effect on the likelihood of systemic breakdown of varying the probability of connection and hence the resulting density of interbank connections (connectivity) is ambiguous and depends on the initial level. For a low initial degree of connectivity, an increase in connectivity unambiguously increases the risk of knock-on defaults. When connectivity is already high, however, a further increase in connectivity tends to help dissipate losses across the system and can thus make the system more resilient to shocks.

The potential for knock-on defaults depends on the aggregate level of capital in the system. The lower the aggregate amount of capital in the system the greater is the potential for knock-on defaults, for any given level of connectivity. Moreover, there is an interaction between the level of capital and the systemic risk created by varying the degree of connectivity. In particular, low aggregate capital increases the potential for interbank connections to transmit shocks through the system. This means that aggregate shocks that might be able to draw down the capital cushion of the system as a whole pose a risk in turning the system from one where connections dissipate shocks to one where connections tend to work as shock transmitters.

The results suggest that when there are few banks, the potential for one failure to entail knock-on defaults is greater, all else equal. This means that more concentrated banking systems are prone to larger systemic risk, all else equal. Further, the risk of contagion (knock-on defaults) is shown to depend on the degree of asymmetry (tiering) inherent in the structure of the banking system. However, the results suggest that tiering does not necessarily lead to greater scope for knock-on defaults. The reason is that the effect of connectivity on the likelihood of systemic breakdown is shown to be ambiguous. When large banks are connected to many other banks this increases the scope for shock transmission. But it could equally lead to better absorption of the initial shock.

Non-linear adjustment of import prices in the European Union

Summary of Working Paper no. 347 José Manuel Campa, José M González Mínguez and María Sebastiá Barriel

The impact that movements in nominal exchange rates have on the geographical allocation of economic activity and the volume of trade has been at the core of research in international economics for over three decades. One key point in this debate is the degree, speed and form in which domestic prices of imported products adjust to exchange rate changes. It is often reported that the high volatility of nominal exchange rates is not matched by the behaviour of import prices, which tend to be far less volatile. This gives rise to fluctuations in real exchange rates (the exchange rate adjusted for relative prices) which have been seen to be large and persistent over the past three decades, suggesting that the adjustment of import prices is very slow.

Several reasons have been suggested for such a slow adjustment of import prices. These include the existence of product differentiation and imperfect competition that can isolate, at least partially, foreign producers' pricing policies from exchange rate changes (implying price differentials between domestically produced and imported tradable products), and the presence of price rigidities driven by some form of fixed cost to changing prices.

Understanding the speed and the form in which the adjustment of import prices — and, thus, real exchange rates — to their long-run equilibrium takes place is an important issue in order to comprehend and anticipate inflation developments and, consequently, to provide an appropriate policy response by monetary policy authorities.

The adjustment of import prices to nominal exchange rate changes has also been an important part of the economic policy debate within the European Union (EU). The adoption of the euro by a subset of twelve countries and the large fluctuations in the value of this currency relative to the US dollar have led to a profound interest in the underlying determinants of import prices and their relationship with exchange rate and monetary conditions.

This paper looks at the process of adjustment of import prices in EU countries towards their long-run equilibrium when they deviate from it due to changes in exchange rates or in foreign prices. The main purpose of the analysis is to gain a better understanding of this adjustment process, in particular by looking at the possibility of a non-linear relationship between deviations from, and adjustments to, the long-run equilibrium (ie there is not a simple proportional relationship between the two). It is possible that prices react proportionally less to small deviations from equilibrium than to large deviations, or that the speed of adjustment back to equilibrium differs when prices are above or below that equilibrium. This is in contrast to the usual assumption that prices adjust linearly; that is, in strict proportion to the size of the deviation. Looking for evidence of non-linearities should help gain a better understanding of this adjustment process. A secondary goal that we try to achieve in this paper is to compare import price adjustment patterns among EU

members that have adopted the euro as their currency and the non euro area countries. If they are different, this could give us some insight into possible structural change when joining a monetary union, which would ultimately affect inflation.

As far as non-linear adjustments are concerned, we considered three different possibilities: that they increase with the size of the deviation (non-proportionality); that they are asymmetric with respect to the sign of the deviation and, finally, that certain thresholds in the size of the deviation exist below which no adjustment takes place. We test these ideas by modelling the process driving foreign prices, nominal exchange rates and import prices in domestic currency allowing for non-linear adjustments. We use a combination of techniques that have been proposed in previous work to estimate such adjustments. We find strong evidence for the presence of non-linearities in the adjustment towards long-run equilibrium in certain industries. This effect is stronger in manufacturing industries. Non-proportional adjustment among manufactures points to the higher degree of price differentiation that characterises these products as an explanation for less adjustment. In contrast, linearity cannot be rejected for agricultural and commodity imports. In some (manufacturing and non-manufacturing) industries, the adjustment is faster the further away current import prices are from their implied long-run equilibrium.

However, in manufacturing there is further evidence of asymmetry in the adjustment to long-run equilibrium: deviations from long-run equilibrium due to exchange rate appreciations of the home currency result in a faster adjustment than those caused by a home currency depreciation. Finally, we also find evidence that prices do not adjust when the deviations are small. We estimate the minimum deviation required for prices to adjust and find that these thresholds tend to be much smaller for manufacturing industries than for commodities.

The resulting evidence points towards adjustment patterns that may differ by country. In general, the patterns of adjustment might be driven by the industry composition of each country's imports and by the competitive structure in each of those industries. In principle, it can be expected that the rate at which cost changes are 'passed through' into prices be lower and less linear in euro-area member states than in countries outside EMU. The reason is that the possibilities for foreign producers to deviate from local producers' pricing policies seem to be less pervasive in larger import destinations. However, the evidence does not point in this direction. In contrast, non euro EU member countries do not appear to have significantly different adjustment patterns from euro-area member states. This suggests that there are no structural differences among these two sets of countries in pass-through rates and that the introduction of the euro, by non euro area member states, is not likely to cause a structural change in this relationship.

The elasticity of substitution: evidence from a UK firm-level data set

Summary of Working Paper no. 348 Sebastian Barnes, Simon Price and María Sebastiá Barriel

The elasticity of substitution (' σ ') between capital and other factors of production, such as labour, is a measure of the ease with which firms can substitute one input for another. For a constant level of output production, if there is a rise in the relative intensity with which one factor is used (eg a rise in the capital to labour ratio), then (except in the extreme case of perfect substitutability, $\sigma = \infty$), firms will have to use increasingly more of that factor (capital) to offset a given reduction in the quantity of the other (labour). At the other extreme, if capital and labour have to be used in rigid, fixed proportions, there is no flexibility, and the elasticity is zero. Generally, σ will take a value between these two limits. This degree of substitution also determines the responsiveness of factor demand to changing prices. Clearly, if factor proportions are fixed, then (for a given output) the demand for capital will not depend at all on the price of capital (or more strictly, the 'user cost' of capital, that takes into account the cost of owning capital, which is affected not only by the price but also by interest rates and other factors). But the larger is σ , the more responsive is capital to changing relative prices. So from the point of view of monetary policy, σ is of interest not least because investment is a major part of the monetary transmission mechanism, responding as it does to changes in interest rates.

Yet there is controversy about its size. Some researchers believe that the elasticity is around unity (the value taken by the widely used 'Cobb-Douglas' production function); others, that it is substantially lower, perhaps below 0.5. To help resolve this uncertainty, in this paper we bring to bear new evidence on investment, output and the user cost from a panel of UK firms.

Much of the debate in the empirical literature is about the proper treatment of short-run dynamics in estimation. The

'around unity' camp argues that there are short-run biases that reduce the estimated value. One of these might follow if firms expect shocks to the user cost to be quickly reversed. They might not then react to such shocks. Some of the US and Canadian evidence using aggregate data seems to support this, although this is not apparent in the UK data. Another potential bias might come from the short to medium-run supply elasticity; movements along a supply curve might be mistaken for shifts along the demand curve, making identification of the demand response (and therefore σ) problematic.

By contrast, some US results from Robert Chirinko and his co-authors using a US panel of firms found a well-determined estimated elasticity of around 0.4. Their method is designed to be unaffected by the dynamic issues mentioned above. By essentially estimating the cross-sectional relationship using 'time averaging' (changes over long periods), it is both simple and robust, and in this paper we apply it to a UK data set. We find similar values to those of Chirinko. Moreover, we find that other methods similarly designed to accurately estimate long-run parameters produce results that are in the same region. This is consistent with results from previous work on aggregate UK data that allow for short-run dynamics.

The main conclusion is that the average elasticity of substitution in our panel of firms is about 0.4, substantially less than unity. This broad conclusion remains largely unchanged no matter which econometric method we use. It is similarly invariant to whether we freely estimate the returns to scale or impose the commonly assumed value of unity (constant returns to scale). So this estimate, consistent with previous work using other UK data sets, seems relatively robust.

Dealing with country diversity: challenges for the IMF credit union model

Summary of Working Paper no. 349 Gregor Irwin, Adrian Penalver, Chris Salmon and Ashley Taylor

This paper assesses the efficiency of the International Monetary Fund's (IMF's) lending framework using a simple, yet novel theoretical model of the IMF as a credit union, in which the membership decides collectively by a vote on the size of the Fund and hence the amount of crisis lending it can provide. This decision, in turn, impacts on individual country choices over the amount of self-insurance to hold in the form of reserves. The equilibrium Fund size and individual country reserve choices are analysed under three different characterisations of the Fund's decision-making processes unconstrained majority voting, constrained majority voting, and qualified majority voting with an agenda setter. The welfare implications in each case are assessed and we consider how the existence of spillovers between countries affects the outcome.

In all cases the analysis suggests the present IMF lending framework may no longer be appropriate. It may well have been during the first two to three decades of the Fund's existence, when almost all countries were potential Fund debtors and had broadly homogenous interests, but the analysis suggests it is much less well suited to the current situation in which members differ sharply in their economic characteristics and needs. In particular, we find that with an increasingly heterogeneous membership, in terms of crisis probabilities, the decisions over the size of the Fund are likely to be driven by members with a relatively low crisis probability. Consequently, the Fund is increasingly unlikely to provide financing on a sufficiently large scale to meet the demands of higher-risk members, leading them to rely more heavily on self-insurance. The analysis suggests that increasing the size of the Fund may be Pareto improving, but only if the financial burden is distributed so that those who benefit most — that is, the countries which have the highest crisis probability — pay the most. This would constitute a significant change in the financing of the Fund's lending operations.

The main message of the paper is that the framework governing the Fund's lending operations may no longer be appropriate. An alternative approach may be needed: one which takes into account that creditor and debtor countries have different interests, but which also takes into account the moral hazard consequences of large-scale lending.

Investigating the structural stability of the Phillips curve relationship

Summary of Working Paper no. 350 Jan J J Groen and Haroon Mumtaz

In the United Kingdom and other G7 countries, the short-run correlation between inflation and different measures of real economic activity has fallen over time and this has coincided with a fall in the level and persistence of inflation. The empirical evidence on shifts in the short-run relationship between inflation and real activity, however, is mainly reduced-form in nature — that is, it cannot by itself tell us about the causal relationships. Therefore, it is not straightforward to use these results to draw inference on the changes in underlying price-setting behaviour of firms. To be able to do that, a structural relationship is needed that embeds the inflation-real activity relationship in price-setting behaviour. We use the New Keynesian structural framework that is popular in the academic literature. This paper is an attempt to assess empirically how the 'deep' parameters that underlie — at least in New Keynesian theory — the price-setting behaviour of firms have changed over time. These 'deep parameters' include the equilibrium rate of inflation, which in the end is set by the central bank, and shifts in that variable can control for shifts in price-setting behaviour that are due to policy regime shifts.

What governs firms' price-setting behaviour? In the New Keynesian framework price stickiness is formalised by assuming that only a randomly selected fraction of firms can optimally reset their prices each period. They do this in the knowledge that their chance to optimise their prices may not be for several periods. Optimal prices are then based on the current level of real costs and inflation expectations. The remaining firms do not re-optimise their prices but instead index their current price increase to last quarter's inflation rate. In this model, the Phillips curve describes how current inflation is affected by current real costs and expected inflation. The contribution of each of these components to current inflation, in turn, depends on the following set of 'deep parameters' that summarise the price-setting behaviour of firms: (a) the fraction of firms that is allowed to freely set their price increase in a profit optimising manner; (b) the average mark-up that firms demand over their costs; (c) the degree of

indexation to lagged inflation for firms that cannot determine their price increase in an optimising manner. A higher degree of indexation is associated with greater inflation persistence; and (d) the level of equilibrium inflation, which determines for the re-optimising firms their perceived risk that future profits can be eroded by increasing inflation. This in turn determines the relative weight the re-optimising firms place on current costs and future risks. In the New Keynesian framework, equilibrium inflation is determined by the central bank.

We estimate the parameters listed above for the euro area, the United Kingdom and the United States, using a method that allows for structural breaks or different states of the world. In other words, the estimated value of these parameters is allowed to be different across subsamples, where the timing and duration of these states is determined endogenously.

What do the key results suggest about the slope of the Phillips curve? As noted earlier, the Phillips curve in this model relates current inflation to current real costs and expected inflation. The weight placed on each component depends on the deep parameters described above. Our empirical analysis indicates that for all three economies only equilibrium inflation shifted over time, whereas the other deep parameters appear to have been unaffected by these shifts. Therefore, shifts in equilibrium inflation have been the main driver in the time variation observed in the slope of the Phillips curve. So our results suggest that the impact of current real costs has increased as equilibrium inflation has fallen. The intuition behind the increase in the impact of current costs is as follows: with low equilibrium inflation (as at present in the United Kingdom under inflation targeting), firms place more weight on current costs when setting prices as future economic conditions, and hence future profits, are more certain. This means that the response to shocks is now more immediate. And as the fall in equilibrium inflation implies that the agents place less weight on future inflation, this will result in a decline in the impact of expected inflation on actual inflation.