
Financial liberalisation and consumers' expenditure: 'FLIB' re-examined

Working Paper no. 157

Emilio Fernandez-Corugedo and Simon Price

The failure of consumption models to predict the fall in the United Kingdom's savings ratio in the late 1980s and its rise in the early 1990s led some economists to look at models of forward-looking consumers who may be unable to borrow. A theoretical weakness in these early papers is that they assumed that the proportion of liquidity-constrained individuals does not change. In the United Kingdom, increased competition in the lending market in the 1980s eased restrictions on borrowers, and is likely to have reduced the number of credit-constrained consumers. So models that assume that the proportion of constrained individuals remains constant through time may not match UK experience. To address this shortcoming, some economists have specified forward-looking consumption functions that assume that the proportion of credit-constrained consumers is inversely related to proxies for financial liberalisation. Their results suggest that this method is able to explain UK consumption data.

This paper examines whether recent UK consumption behaviour can indeed be explained using this method. To this effect we construct a proxy for financial liberalisation (FLIB). FLIB is defined as the sum of the constant and the residuals in a regression of the loan to value ratio on the house price to income ratio, the nominal post-tax mortgage rate and a two-year moving average of the post-tax mortgage rate. We then re-examine a forward-looking consumption model which uses FLIB as a variable identifying the proportion of liquidity-constrained individuals. We find that this

implementation of the model, which inevitably embodies joint hypotheses about consumption behaviour and about the measurement of financial liberalisation, is not robust and does not give a plausible picture of the number of people who were liquidity constrained in the 1990s.

We argue that one possible explanation for these results is that the liberalisation proxy is unable to depict accurately the consequences of UK financial deregulation in the 1990s. FLIB's behaviour in the 1990s suggests that all the liberalisation that occurred in the 1980s was reversed the following decade, which seems implausible. This in turn means that, by assuming that the proportion of constrained agents in the economy is a function of FLIB, the consumption model examined in this paper does not derive a plausible measure of this key variable. We argue that the mapping from the FLIB index of liberalisation to the proportion of constrained consumers is somewhat arbitrary and that some of the assumptions made to derive a functional form for an estimatable UK consumption function might account for the failures encountered in this paper.

Finally, we attempt to explain the behaviour of FLIB in the 1990s. We identify the sharp reduction in the nominal interest rate as the main factor accounting for FLIB's reversal over the 1990s. We also argue that lending institutions may have changed the emphasis of their lending criteria towards loan to income ratios after liberalisation.

Soft liquidity constraints and precautionary saving

Working Paper no. 158

Emilio Fernandez-Corugedo

Miles Kimball defines a precautionary motive as ‘any aspect of an agent’s preferences which causes a risk to affect decisions other than the decision of how strenuously to avoid the risk itself and risks correlated with it (which is governed by risk aversion). A precautionary motive leads an agent to respond to a risk by making adjustments that will help to reduce the expected cost of the risk’. Thus, precautionary saving arises when forward-looking consumers accumulate wealth today for the purpose of reducing the impact of future uncertainty on future consumption decisions. Liquidity constraints arise when consumers have difficulties to obtain credit. More specifically, soft liquidity constraints represent the situation where consumers are able to borrow, but incur penalties which increase with the amount borrowed. Hard liquidity constraints refer to the unavailability of credit altogether.

The modern consumption literature has examined the problem of how much to consume and save each period under two polar scenarios. One scenario considers perfect capital markets where no barriers to borrowing exist and where interest rates are the same for savers and borrowers. The other scenario assumes that consumers are not able to borrow at all. Both scenarios, however, do not seem to match what is commonly observed in developed economies: consumers often borrow and face interest rates that are higher for debt than for saving.

The theoretical implications for consumption arising from the two polar cases are summarised by Carroll and Kimball in two papers that provide the conditions under which the introduction of uncertainty and liquidity constraints leads to precautionary saving, and analyse how precautionary saving. Technically, these conditions require the interaction of risk (either to labour income or to the rate of return) with liquidity constraints and/or with certain functional forms for the utility function.

The literature finds at least three important implications of the inability to borrow (hard constraints) for consumption. First, hard constraints increase precautionary saving around levels of wealth where the constraints bind. Second, if consumers face the possibility of becoming constrained at any point in the future, they will behave as if they were constrained today, even in the absence of a current liquidity constraint. Finally, the introduction of further borrowing constraints does not necessarily lead to an increase in precautionary saving.

This paper considers the implications for consumption behaviour when households are allowed to borrow, but face penalties that increase with the amount borrowed. The introduction of this type of constraint does not lead to consumers behaving very differently from consumers who face hard constraints. A soft constraint increases precautionary saving and affects prior periods, although the introduction of further soft constraints can lead to lower precautionary saving. However, a new result is that the amount of precautionary saving is reduced when hard constraints are relaxed and become soft. The intuition behind this result is simple: when consumers cannot borrow, they must have savings to avoid shocks that could leave them with low levels of income. A relaxation of the borrowing constraint means that consumers do not need to have these (high) savings to avoid adverse shocks to income. More technically, the paper shows the effects that soft liquidity constraints have on the value, marginal value and consumption functions in a dynamic programme. The introduction of a soft constraint makes consumers more averse to risk (since the value function becomes more concave) and also more prudent (since the marginal value function becomes more convex). An implication is that the resulting consumption function becomes concave with respect to wealth.

The implications of an ageing population for the UK economy

Working Paper no. 159

Garry Young

This paper discusses the impact of demographic change on the UK economy, looking at effects on GDP growth and GDP per head, saving and capital investment, interest rates, asset prices and the distribution of national income. It also considers the risks associated with demographic change. A key finding, widely supported in the academic literature, is that even under relatively cautious assumptions about technological progress and capital accumulation, aggregate living standards (as measured by GDP per head) are set to double over the next 50 years. While there are clear risks to this aggregate outlook, these would be present even without demographic change.

The impact of ageing on the rate of saving and capital accumulation is one of the key uncertainties surrounding any projection of long-term growth. The paper analyses this in the context of a model where people are reliant on their own saving for their retirement income and considers three different types of demographic shocks: a baby boom, an increase in longevity and a decline in fertility. The overlapping generations model used for this purpose makes it possible to assess the impact of these shocks on the welfare of different generations under different assumptions about household behaviour. It finds that a baby boom has an adverse effect on the baby boom generation for the obvious reason that when they are of working age their abundance drives down wages and when they are of retirement age the abundance of their saving drives down the rate of return. The impact of a baby boom on other generations is largely beneficial. Increases in longevity, not accompanied by changes in labour supply, have a detrimental effect on annual consumption per head for the obvious reason that people have more years over which to spread their consumption. Changes in fertility appear to have very little effect on individual consumption per head, although they clearly affect aggregate quantities because of changes in the number of people.

An important conclusion of these models is that while individual consumption over the life cycle may not be strongly affected by demographic change, there can be large effects at particular parts of the life cycle when individuals do not attempt to spread their consumption evenly. For example, the analysis of greater longevity suggests that this might reduce individual life-time consumption by about 2% if the change is spread evenly over time. But if individuals follow rule-of-thumb behaviour prior to retirement and do not accumulate enough assets, the reduction in their life-time consumption will be concentrated into the years when they are old. This is particularly important at the current juncture

since many people in prime saving age will observe their own pensioner parents living longer without any obvious adverse effect on their consumption. This could be misunderstood as suggesting that their own saving for retirement is adequate. Yet the formal model suggests that the early generations to benefit from greater longevity do not have to reduce their consumption since the capital accumulated by previous generations is not affected by their longevity. But their children will receive smaller bequests. Moreover, the current generation of pensioners has benefited from extraordinarily high asset returns which are unlikely to be repeated.

The implications of this analysis for interest rates are modest. This is consistent with other research which suggests that the effect of demographic change on asset prices more generally is likely to be small. This leads on to the second conclusion of this paper, that the risks to the living standards of individuals and individual cohorts are large. While the impact of demographic change on asset prices is small, the historical volatility of asset prices and rates of return is significant. This is unlikely to be affected by demographic change, but it means that those relying on financial market returns for their retirement income could be much less lucky than those who enjoyed the high returns of the 1980s and 1990s.

Moreover, the projected increase in the number of people in this position raises the risks of large numbers suffering the effects of financial shocks, as well as the risks to macroeconomic and financial stability. Recent experience with endowment mortgages emphasises that the returns on long-term investments can turn out to be substantially different to expectations. In a similar way, a period of very low rates of return on capital would leave people with much lower pension entitlements than had been anticipated. This can occur even when overall asset returns have been strong if investors have poorly diversified portfolios, but the adverse effect of it occurring for a substantial group of savers could be severe. Such an outcome would have macroeconomic repercussions if lower expenditure by the retired was intensified by lower spending by those of working age who become concerned about their own retirement income. It would have systemic implications if lower asset returns meant that debts could not be paid.

Given the lack of financial sophistication of many households, there is a clear educational role for financial regulators in informing people of the risks they face and what action they might take.

On gross worker flows in the United Kingdom: evidence from the Labour Force Survey

Working Paper no. 160

Brian Bell and James Smith

Empirical studies of worker flows in the United States and Europe have found that these flows are large when compared with the change in the stocks of employment and non-employment and have a distinct cyclical pattern. In the United Kingdom, studies of this kind have been hampered by limitations in the available data. In this paper we make use of newly released longitudinal data from the Labour Force Survey (LFS) to document the size and cyclical patterns of the gross worker flows in the United Kingdom.

The motivation for considering gross worker flows is a simple one: to uncover what lies behind the headline levels of—and changes in—key statistics such as employment and unemployment. In particular, data on gross worker flows allow us to observe two features of these flows: their magnitude and cyclical properties. The magnitude of worker flows may allow us to gauge the flexibility of an economy, as the rate at which workers flow from less efficient plants to more efficient ones will affect how quickly an economy responds to economic shocks. And the cyclical properties of the gross flows allow us to uncover how labour demand is met over the business cycle. In short, the availability of data on gross worker flows allows us to go behind the aggregate stock data to examine the nature of labour market dynamics.

Data on gross flows may be affected by measurement biases to a greater extent than the levels data. In particular, sample attrition and response error may cause errors in estimating the flows. We test this by looking at the number of ‘inconsistent’ transitions. In the LFS, individuals in employment and unemployment are asked not only about their current state, but also how long they have been in that state. If the duration contradicts the transition, then the transition is ‘inconsistent’. We observe a significant level of inconsistent transitions, but suspect that most of the error occurs because individuals are unclear as to their exact duration in any state rather than about their current state. To the extent that these transitions are not genuine, they will lead to overestimation of the gross flows.

Over the past five years, the stock of unemployed fell by an average of 40,000 per quarter. Given an average stock of 1.9 million, this may seem to suggest that the market for labour can be characterised as fairly static. Yet such a conclusion would be wrong. We find that, over the same period, almost three-quarters of a million people entered unemployment in a quarter, with numbers drawn equally from employment and inactivity. Similarly, almost one million

people start a new job each quarter after previously being unemployed or inactive.

Theoretical models of labour market flows generate predictions about the cyclical pattern of flows and associated hazard rates (the chances of making a transition from a given labour market state to another). These predictions can be tested using the LFS longitudinal data. In particular, we examine the cyclicity of both the gross flows and the associated hazard rates in the United Kingdom using a variety of data and techniques. We find that:

1. Flows from employment to unemployment are countercyclical, as is the hazard rate. The reverse flow, from unemployment to employment, is also countercyclical—while its associated hazard is strongly procyclical.
2. Flows from employment to inactivity tend to be procyclical and there is no clear pattern to the associated hazard rate. Flows from unemployment to inactivity appear to be countercyclical.
3. Flows and hazards from inactivity are imprecisely measured, and we cannot be confident of any statement on their cyclical characteristics.
4. Flows of workers moving from one job to another, without a recorded period of unemployment or inactivity, are strongly procyclical.

These findings are broadly consistent with similar results for the United States and Europe.

In addition, we are also able to measure the incidence of job-to-job flows. Little is known about these flows in the United Kingdom and previous research has tended to focus on the prevalence of on-the-job search without knowing whether that search was successful. We show that 2.9% of those in employment change employer in an average quarter. This represents a movement of three-quarters of a million workers. Unsurprisingly, the probability of making such a move is much higher for those who are engaged in on-the-job search. Such movements tend to occur much more frequently for workers with short tenure in their initial job. This is consistent with findings in the literature suggesting that individuals search on the job when they are in poor matches. As tenure lengthens and job-specific human capital is acquired, the incentive to move jobs falls.

Regulatory and ‘economic’ solvency standards for internationally active banks

Working Paper no. 161

Patricia Jackson, William Perruadin and Victoria Saporta

The Basel Committee is currently engaged in designing a new Accord on the capital adequacy of internationally active banks that will supersede the original 1988 Accord. Under the old Accord internationally active banks in G10 countries are required to hold broad and narrow capital that is no less than 8% and 4% respectively of their risk-weighted assets. While exposures to banks, sovereigns and mortgage assets are treated differently, the bulk of banks’ private sector assets are subject to the same capital charges no matter how risky they are. In framing the new Accord, the Basel Committee intends to make capital charges on different exposures much more risk sensitive.

A major question confronting regulators is how high the overall average level of capital charges for representative banks should be. The 8% in the original Basel Accord was chosen on the basis that this was the minimum level of capital observed among banks that were perceived to be following best industry practice. When capital charges are being redesigned as part of the preparation of the new Accord, it is natural to consider what current levels of regulatory capital imply for financial stability and to what extent these are a binding constraint on banks.

This paper employs a standard credit risk model to investigate the survival probabilities implied by the current minimum level of narrow regulatory capital (subordinated debt included in the wider definition does not affect survival probabilities). In particular, for corporate loan portfolios with representative quality distributions, we calculate the likelihood that the banks holding the portfolio would survive over a one-year horizon if their capital were at the regulatory minimum. The model employed is a simplified version of the widely used CreditMetrics approach, developed by JP Morgan.

We then compare the survival probabilities implied by minimum levels of regulatory capital with those levels of ‘economic’ capital that internationally active banks actually hold. We do this in two ways. First, we again employ CreditMetrics calculations. Second, we examine the historical survival rates associated with the ratings that banks receive from the main rating agencies. This latter examination is complicated by the fact that, in many cases, banks’ agency ratings are boosted by market expectations that the authorities will provide support if banks experience difficulties. To get around this

problem, we calculate, using an econometric model, what ratings banks would have if they were not expected to be able to obtain support.

We conclude from our calculations that the one-year survival probability or solvency standard implied by the current Basel Accord minimum capital levels is between 99.0% and 99.9% depending on the quality of the corporate loan book used. Our investigation of the solvency standard implicit in the ‘economic’ capital levels that internationally active banks actually hold suggests that it is substantially higher than 99.9%.

Having investigated the relation between the survival rates implied by regulatory minimum capital levels and by the capital levels that banks actually hold, we ask a further question: why do banks make such apparently conservative capital decisions, selecting economic capital that significantly exceeds the regulatory minimum? It is difficult to answer this question conclusively but we argue that the evidence is at least consistent with one explanation, namely that banks are obliged to maintain higher capital levels in order to obtain access to certain wholesale markets, most notably the swap market, participation in which is a prerequisite for operating a modern large-scale, internationally active bank. To make the case that market discipline of this kind is an influence on banks’ choices, we show that the volume of banks’ swap liabilities, conditioning on bank size, is significantly correlated with the bank’s credit rating. Large international banks wanting to deal in significant swap volumes appear to have to maintain high ratings.

The main implication of our analysis is that maintaining minimum regulatory capital levels in the new Basel Accord at levels similar to those that apply under the 1988 Accord would not act as a major constraint on most internationally active banks, since they already operate on higher ‘economic’ solvency standards than those implicit in the Basel regulatory minimum. While different reasons might be adduced for why banks adopt a relatively conservative approach in their capital-setting decisions, one possibility that seems consistent with data on swap market volumes is that the need to maintain access to certain wholesale markets, which is crucial to operating a large bank, necessitates a fairly stringent ‘economic’ solvency standard.

Factor utilisation and productivity estimates for the United Kingdom

Working Paper no. 162

Jens Larsen, Katharine Neiss and Fergal Shortall

Capacity utilisation—or time-varying factor input utilisation—is a key component of the supply side of the economy and is often thought to provide information regarding the build-up of inflationary pressures. Though difficult to measure, capacity utilisation may also account for much of the variation in aggregate output. So it may provide useful insights into the characterisation of the business cycle.

This paper evaluates the implications of a general equilibrium model of time-varying factor utilisation under the assumption of *factor hoarding*. We assume that production relies on labour and capital services. The contribution from labour depends not only on the number of hours people work, but also on the productive effort they exert during those hours. Similarly, the contribution from capital takes into account not only the number of machines in the factory, but also the intensity at which they operate. There are costs associated with utilising factors intensively: workers suffer disutility, machines wear out more quickly. But since firms must choose in advance how much capital stock and employment to rent, they may under-utilise these inputs in equilibrium. Machines may be left idle; workers may spend time sweeping the factory floor. We find that firms initially respond to unanticipated shocks by altering factor utilisation rates. In subsequent periods, firms adjust their physical stock of capital and employment. As a result, utilisation rates are a leading indicator of firms' hiring of both capital and labour.

We then use the model to derive estimates of capital utilisation and labour effort for the United Kingdom. By explicitly accounting for variations in factor utilisation, these help to estimate total factor productivity (TFP)—that portion of output growth not due to growth in capital or labour more accurately.

Our estimate of capital utilisation for the United Kingdom matches survey-based measures of capacity utilisation quite closely, supporting the view that these measures accurately reflect the degree to which firms are utilising their existing capital stock. All measures

indicate that capital utilisation rose during the 1990s—though it has recently fallen back somewhat—reflecting a declining capital-to-output ratio over the period. The predicted positive and leading relationship between capital utilisation and investment in the model in turn indicates the potential usefulness of surveys for forecasting investment.

Movements in total hours worked drive our estimate of labour effort. Given the costs to adjusting employment, this is quite intuitive. When a boom is in its initial stages, firms demand an increase in effort in order to generate labour services. Only after a time can firms satisfy their demand for increased labour services by increasing total hours worked, with effort slowly returning to normal levels. Our estimated series for labour effort shows a decline after the mid-1990s. This decline is a reflection of the sharp increase in total hours worked over that period. Contrary to theoretical predictions, however, our effort series is only weakly correlated with both a manufacturing-based measure of labour effort and average hours worked.

Our estimate of TFP is found to be less cyclical than the traditional measure, the Solow residual. Nevertheless, a weighted average of capital utilisation and labour effort—which we call aggregate factor utilisation—is not closely related to the Solow residual. This suggests that measures that conflate both capacity utilisation and temporary fluctuations in TFP (as the Solow residual does) may be misleading indicators of excess demand pressure.

Rather, our measure of aggregate factor utilisation is more correlated with detrended labour productivity. In some ways this is not surprising: if capital and labour are slow to adjust, then much of the variation in factor inputs—and hence output—over the business cycle must come from utilisation and effort. This supports the view that labour hoarding is responsible for much of the cyclicity in measured labour productivity. In fact, labour productivity, when calculated as output per unit of *effective* labour input, is much less cyclical than a simple measure of output per hour.

Productivity versus welfare: or, GDP versus Weitzman's NDP

Working Paper no. 163

Nicholas Oulton

Traditionally, productivity at the aggregate level has been measured using GDP, ie the measure of output is gross of depreciation. But suppose that the composition of the capital stock is shifting towards assets with shorter lives, so that the average depreciation rate is rising. This suggests that some part of what GDP measures as an increase in output may be illusory: some of the extra output is needed just to maintain the capital stock at its existing level. This question is given a sharper focus by the experience of the United States in the 1990s, where the growth rates of labour productivity and total factor productivity (TFP) rose, while investment shifted towards short-lived ICT assets. This raises the possibility that the US productivity improvement might be just a statistical illusion.

This paper has a theoretical and an empirical part. In the theoretical part, I compare measures of productivity with measures of welfare. I conclude that while GDP is satisfactory as a measure of *output*, it is outclassed as a measure of *welfare* by what I call Weitzman's NDP (WNDP). This is nominal net domestic product (consumption plus *net* investment) deflated by the price index for consumption. I extend the theory behind WNDP to the case where TFP growth can vary across sectors. This is the empirically relevant case for analysing recent US experience.

The aggregate TFP growth rate is the rate at which the GDP frontier is shifting out over time. This can be decomposed into a weighted average of the TFP growth rates in the various industries. Analogously, we can define the rate at which the WNDP frontier is shifting out over time. I call this the growth rate of total factor welfare (TFW). Like aggregate TFP growth, TFW growth can also be decomposed into a weighted average of TFP growth rates in the various industries, but the weights are not the same as for the GDP frontier. Hence the growth of welfare over time can be analysed using the

same tools as have been developed for the analysis of the growth of output.

In the empirical part of the paper, I apply some of these ideas to the experience of the United States in the 1990s. In principle, one might expect WNDP to have grown more slowly than GDP over this period, for several reasons. First, the weight on consumption is higher in WNDP (or in NDP) than in GDP and consumption has been growing more slowly than investment. Second, the relative price of investment goods has been falling and this reduces WNDP growth. Third, one might have expected depreciation to have risen as a proportion of GDP, thus raising the share of consumption in WNDP still further.

In practice, WNDP has grown a bit more slowly than GDP. But the gap between the two growth rates was actually somewhat larger in the period 1973–90 than it was post 1990. And the acceleration of WNDP post 1995 was equal to that of GDP. The explanation is twofold. The ratio of depreciation to GDP has in fact been stable, despite the growing importance of short-lived assets. And net investment has grown more rapidly than gross investment. The growth rates of TFP and of TFW in the US non-farm business sector are also compared and found to be similar in the 1990s. Moreover, they display an almost identical increase after 1995.

GDP is a measure of output, not of welfare. So even if GDP had grown significantly faster than WNDP, this would not by itself suggest measurement error. In fact, the two have grown at similar rates in the 1990s and accelerated by the same amount. So it seems that, in practice, GDP has provided as reliable a measure of the improvement in US living standards over this period as WNDP, even though WNDP is conceptually superior as a welfare measure.

Understanding UK inflation: the role of openness

Working Paper no. 164

Ravi Balakrishnan and J David López-Salido

This paper re-examines inflation dynamics in the United Kingdom. Our main motivation is the recent low inflation, low unemployment era in the United States, the United Kingdom and the euro area. This has led to overpredictions of inflation using standard specifications of traditional Phillips curves. This has been a major motivation for a ‘New Phillips Curve’ approach, which has had success for the United States and the euro area. The reason the United Kingdom is an interesting case to study is that it is a far more open economy than the United States or the euro area.

In this paper we analyse whether the openness can explain the overprediction problem in traditional Phillips curve estimates and whether it affects the performance of ‘New Phillips Curve’ estimates. The paper is divided into two parts. First, we document the overprediction problem for the United Kingdom and try to solve it in a traditional Phillips curve framework. We introduce external shocks from two sources: terms of trade shocks and domestically generated inflation (DGI). We find that external shocks do not fully solve the overprediction problem within this framework. We further argue that there is a more general misspecification problem with traditional Phillips curve estimates, due to the presence of regime changes and structural change in the UK economy.

Second, we look at ‘New Phillips Curve’ estimates. They do not perform particularly well: real marginal cost is not significant in our baseline specification. Further investigation suggests the relationship between marginal cost and inflation broke down around the mid-1980s. When we use a labour-share measure adjusted for the public sector, real marginal cost becomes significant, but the goodness of fit of the model—based on fundamental inflation—is still very poor.

Next, we extend our ‘New Phillips Curve’ model to allow for open-economy influences. In particular, we take into account imported intermediate goods. When we allow for imported intermediate goods the relationship between inflation and marginal cost improves significantly. Fundamental inflation performs better than previously, but still has a tendency to underpredict

and then overpredict inflation: something also present in the traditional Phillips curve estimates.

Finally, we decompose the open-economy measure of marginal cost to learn more about its driving forces. We find that a wage mark-up component is important and highly countercyclical. We also find that relative price movements, of taxes relative to overall prices and of imported intermediate goods relative to wages, have been a negative influence on marginal costs over the 1990s. Understanding likely future developments in these relative prices could contribute to the assessment of prospects for marginal costs and the pressures on inflation.

Time-varying desired mark-ups may, in part, explain why the open-economy New Phillips Curve still underpredicts and then overpredicts inflation. In the models considered in this paper, the desired mark-up is assumed to be constant. This is important to the extent that the desired mark-up varies cyclically and can be influenced by external factors. For example, recently there has been much speculation that the high level of sterling has forced manufacturers to cut their margins on exported goods. This is equivalent to a fall in the desired mark-up and will have a negative impact on inflation in the GDP deflator. This idea fits well in a customer market model. In a customer market model, firms are assumed to be monopolistically competitive, and set their own mark-up, taking the mark-up of other firms as given. However, there is a dynamic element to the firm’s problem in that higher relative prices reduce market share. In addition, some consumers are assumed to pay a cost when switching from one firm to another. This kind of model provides a justification for firms to allow the desired mark-up to vary, in the short term, in order to stop the long-term loss to profitability of losing customers. It may also be a key factor that exporters take into account, by allowing margins to vary in reaction to changes in exchange rates, rather than the foreign price of the exported good. Recent high levels of sterling may have reduced the desired mark-up and thus potentially explain the overprediction of actual inflation by fundamental inflation. We plan to look at this in greater detail in future work.

Committees versus individuals: an experimental analysis of monetary policy decision-making

Working Paper no. 165

Clare Lombardelli, James Proudman and James Talbot

Evidence from around the world suggests that the majority of central banks take monetary policy decisions by committee rather than through a single individual. Despite this observation, there is little direct empirical or theoretical evidence on the relative merits of monetary policy decision-making by committees versus individuals. Recent work by Blinder and Morgan has sought to shed light on this question by taking an experimental approach, the main result being that the decisions of committees were superior to those of individuals. Although the results of our paper support this conclusion, we attempt to extend their work by examining and testing several hypotheses as to why this improvement might come about.

To this end, we asked a large sample of economically literate undergraduate and postgraduate students from the London School of Economics to play a simple monetary policy game. Participants acted as monetary policy makers, setting interest rates to ‘control’ a simple macroeconomic model calibrated to match UK data and subject to an unknown combination of shocks. Each participant acted as both individual decision-maker and as part of a committee of five players. All players faced an identical incentive structure: performance was judged according to a score function that penalises deviations of output and inflation from their target values; and they were paid according to their performance.

Just like actual policy-makers, participants in our experiment were forced to make decisions in an uncertain world, while observing only the evolution of the endogenous variables over time. As in real life, these monetary policy makers did not know with certainty the exact structure of the economy they were attempting to analyse. To the extent that players came to the experiment with different prior beliefs about the structure of the model, they may have responded differently to the same set of shocks. So we modelled

these differences of opinion by asking participants to fill in a questionnaire that attempted to reveal these prior beliefs. By asking players to fill in the same questionnaire at the end of the game we were able to discern some evidence of players learning about the underlying model of the economy over time. And for the ‘worst’ players, their improvement in scores over time was positively and significantly related to the extent of their learning about the underlying model.

Like Blinder and Morgan, we found that committees performed significantly better than the individuals who composed them. There are several competing hypotheses as to why. Our results suggest two reasons why committees make better decisions. First, collective decision-making appears to give more weight to the better and less weight to the worse committee members—as judged by their scores when playing the game as individuals—than would be implied by taking the mean of their individual performance. But we find evidence that committees do more than this, enabling all members to improve their performance by sharing information and learning from each other. For example, the performance of the committee was on average better than that of its ‘best’ policy-maker when playing alone.

In our experiment, we also explicitly tested whether the ability to discuss a decision drives the observed improvement in performance. In practice, this did not appear to be the case: in our simple monetary policy game, participants were able to share enough information by simply observing each other’s behaviour. But we were able to illustrate how the relative importance of different types of communication depends upon the nature of the decision problem in a variant of the game in which we slightly altered the structure so as to raise the relative importance of discussion. When we did so, committees that discuss performed better.

The role of corporate balance sheets and bank lending policies in a financial accelerator framework

Working Paper no. 166

Simon Hall and Anne Vila Wetherilt

Recent financial crises have illustrated that the financial positions of borrowers and lenders—financial stability considerations—can influence the way in which official interest rate changes affect spending and inflation—monetary stability considerations. A substantial academic literature has developed considering potential macroeconomic impacts of financing decisions by borrowers and lenders. Among these so-called ‘credit channel’ models, the recent financial accelerator approach of Bernanke, Gertler and Gilchrist seems particularly suited to an analysis of how corporate sector balance sheets and the behaviour of banks can affect the monetary transmission mechanism.

In credit channel models, firms often find it more costly to finance investment projects with external funds rather than with internally generated resources. This ‘*external finance premium*’ may arise because lenders face costs from observing and/or controlling the risks involved in supplying funds to borrowers. These agency costs, and the external finance premium, may vary with borrowers’ financial health. For example, the stake of a borrower in an investment project (measured by the degree to which it is able to finance a project using internal funds) may provide a signal of the unobserved risk of lending. It may also affect the borrower’s incentive to act diligently and to report project outcomes truthfully.

The financial accelerator model used in this paper embeds a similar imperfect information problem in the supply of external finance in a standard macroeconomic framework. Our paper examines how a range of interesting financial scenarios can arise out of this model and in turn, how these scenarios affect the dynamic response of the model economy to alternative shocks (for example monetary news or productivity shocks). These scenarios are defined in terms of steady-state credit spreads, bank lending policies and corporate financial health. The main objective is to examine how the strength of the monetary transmission mechanism might vary across such scenarios.

Our simulations of the model show how balance sheet positions of the financial and non-financial corporate

sectors can affect the monetary transmission mechanism. We show that in certain financial scenarios the financial accelerator mechanism is very potent, whereas in others it has little incremental impact. This implies that, for a given shock in the model economy, monetary policy can be less or more proactive, respectively.

In addition, the model simulation results suggest that certain parameters may merit particular attention. For example, the sensitivity of bank lending to news about corporate financial health has an especially marked impact on the model’s dynamics. And as illustrated in previous work, leverage also plays an important role in amplifying and propagating shocks. But we also show that the strength of the financial accelerator cannot be attributed to a single variable. For example, we observe that the financial accelerator can be weak, both when leverage is low and banks are relatively restrictive in their lending, and when leverage is high and banks are very accommodative.

These theoretical results are consistent with real-world experience that bank and non-financial corporate balance sheets can, at times, have a marked impact on the effectiveness of monetary policy. But while the specific model used in this paper provides an attractive analytical framework for thinking about potential qualitative effects of changes in financial conditions on real variables, we think its quantitative results need to be interpreted with caution, as in all calibrated simulated models. Moreover, although the model can be used to analyse certain important interactions between financial imperfections and the monetary transmission mechanism, it leaves out several features that one might want to incorporate in a more general model of financial stability. For example, the model has a relatively restricted financial structure with a focus on debt finance. Financial institutions are sparsely modelled, with limited potential for effects from the bank lending channel. That suggests that further work to develop quantitative models that incorporate these features may provide further insights into interactions between monetary and financial stability.