Wealth and consumption: an assessment of the international evidence

Working Paper no. 275

Vincent Labhard, Gabriel Sterne and Chris Young

Since the mid-1990s, there have been remarkable changes in stock market capitalisation in many of the major economies, related in part to changes in the valuation of equities. Between Autumn 1994 and Autumn 2000, stock market capitalisation increased by over 100% of GDP in the United Kingdom and the United States; while between 2000 Q3 and 2003 Q2, market capitalisation fell by 65% of GDP in the United States and by 87% in the United Kingdom. These changes have motivated renewed interest in the wealth effect on consumption.

The wealth effect on consumption is often captured by the marginal propensity to consume from financial wealth (mpcw). The existing empirical literature suggests that this quantity varies greatly across countries, and new results presented here, based on single-country structural vector autoregressions (VARs) for eleven OECD countries, tend to confirm this finding. This divergence though is at odds with the values used in calibrated models, which tend to be more similar across countries. The main objective of this paper is to offer a critique of the literature, and to assess several possible explanations that might justify such differences in the mpcw across countries.

It is concluded that many potential explanations cannot account for the magnitude of the differences reported in the empirical literature on the mpcw, including differences in demographics and in the type of assets held across economies. It is argued that there is little theoretical rationale for such a wide cross-country dispersion of empirical estimates. In part, this may be due to the empirical approach taken in much of the literature: partial equilibrium approaches to capturing the impact of changes in wealth on consumption face a cocktail of data problems and cannot account for underlying structural causes of simultaneous changes in both consumption and wealth. For example, in circumstances in which there are shocks to expected earnings, economies where market capitalisation is low and wealth held in unquoted equities is underrecorded might be (inaccurately) estimated to have a higher mpcw. Because conventional empirical estimates of the mpcw are commonly calculated by dividing an empirical estimate of the partial elasticity of consumption with respect to wealth by the observed wealth to consumption ratio, if the reaction of consumption to an earnings shocks is similar, but wealth is underrecorded because of data problems, then the mpcw will be overestimated.

This leaves the question of how it may be possible to assess empirically the likely wealth effect on consumption. Using a suitable panel technique we find that the hypothesis of the long-run mpcw being the same across countries cannot be consistently rejected, and obtain a plausible estimate for the cross-section of eleven OECD countries. This estimate is a little over 6%, broadly consistent with estimates used in a wide range of policy models.

Corporate expenditures and pension contributions: evidence from UK company accounts

Working Paper no. 276

Philip Bunn and Kamakshya Trivedi

Understanding how companies react to financial pressures is an important academic and policy concern. Apart from being relevant to any comprehensive appreciation of corporate behaviour, the ways in which companies adjust their balance sheets and the size of those responses are also inextricably linked to financial stability risks. In this paper we focus on one specific source of financial pressure — contributions to company pension schemes — and investigate the empirical relationship between corporate expenditures and variation in pension contributions within a panel of non-financial UK firms.

Contributions to fund shortfalls in defined benefit pension schemes are a useful example of financial pressure because these must often be made by the sponsoring companies in line with regulatory requirements, and therefore constitute a relatively exogenous source of variation in internally generated finance. The presence of a budget constraint implies that such contributions to the pension scheme divert cash from alternative uses such as dividend payouts or investment. If a firm is financially constrained, or if external finance is costly, pension contributions could force a company to cut dividends and/or not make investments it might otherwise have pursued. Indeed, if the funding positions of pension schemes are related to the stock market (say because they are all exposed to similar equity markets) then financial pressures may affect many companies at the same time, and individual company level responses may add up to large systemic effects.

The recent financial difficulties of company pension schemes resulting from falling asset values and declining interest rates have been highlighted by the introduction of new accounting standards such as FRS17 in the United Kingdom. There have been many estimates of the magnitude of these difficulties, but they all imply that an economically significant increase in pension contributions would be required by sponsoring companies to eliminate the current deficits faced by defined benefit pension schemes in the absence of a sustained rise in the stock market. The results of this paper can therefore be used to inform an assessment of the possible implications of these increases in contributions for company balance sheets.

Using a panel of quoted non-financial UK companies between 1983 and 2002 we estimate generalised methods of moments models for dividends and investment based on those in the existing literature. The main innovation of this paper is to augment these models with a measure of company pension contributions. Our results suggest that firms pay lower dividends than they would have otherwise done in response to increases in pension contributions, controlling for other components of the balance sheet such as capital, cash flow, debt and investment. Dividends have a similar elasticity with respect to pension contributions as cash flow; this is plausible because pension contributions are effectively reductions in corporate cash flow. But this marginal effect implies that dividends are not reduced in response to higher pension contributions on a one for one basis, and therefore in the presence of a binding budget constraint there may also be further adjustment through other channels such as investment, debt or equity issuance. Empirically we find only weak evidence that firms reduce their investment in a statistically significant way as pension contributions rise. This result implies that adjustment to company balance sheets on account of increased financial pressure from higher pension contributions comes mainly through financial rather than real channels.

When is mortgage indebtedness a financial burden to British households? A dynamic probit approach

Working Paper no. 277

Orla May and Merxe Tudela

Since the mid-1990s the volume of secured lending to households has expanded rapidly, both in absolute terms and in relation to household incomes. In 2004, the stock of secured lending to households exceeded £850 billion (compared to around £400 billion in 1995) and represents the largest domestic on balance sheet exposure of UK-owned banks. The rates of arrears and write-offs on secured debt have fallen in recent years and, despite a slight pickup in the second half of 2004, are currently at historically low levels. But there is a risk these could rise further if households began to encounter problems servicing their mortgage debt.

This paper seeks to explain the determinants of mortgage payment problems using disaggregated data from the British Household Panel Survey (BHPS). By using disaggregated data, we can examine how both macroeconomic factors (such as interest rates and house prices) and household-level factors (such as employment status and saving behaviour) affect the probability of households meeting their mortgage commitments. Since the BHPS is a panel survey, it allows us to track the same individuals over time; so we can also examine the dynamics of mortgage payment problems. In particular, we can analyse whether changes in a person's circumstances (such as changes in income) and previous experience of payment problems affect their current ability to service mortgage debt.

The data confirm that the two most important household-level factors associated with mortgage payment problems are adverse changes in employment and the level of income gearing (the ratio of mortgage payments to household income). Becoming unemployed significantly increase the probability of mortgage payment problems. But the results show that if the household is persistently unemployed this is not associated with a higher probability of payment problems, presumably because the household can adjust consumption so that servicing the mortgage is no longer a problem. However, this result may be driven by the fact that there are only a small number of mortgagors in our sample who are unemployed for two or more years.

We find evidence of a positive relationship between income gearing and the probability of mortgage payment problems — a higher level of income gearing significantly increases the probability of payment problems. However, this relationship is only apparent when gearing passes 20% — below that level there is no significant effect on payment problems from income gearing.

The level of effective mortgage interest rates is also found to increase the probability of mortgage payment problems. This is the only non-household-specific variable that is found to have a significant effect. The aggregate level of unemployment has no independent effect beyond that identified at the household level.

The results also show that problems paying for secured debt are persistent. The experience of payment problems has a genuine behavioural effect upon the household in the sense that previous experience of problems increases the probability that the household will subsequently have difficulty servicing its mortgage. There are a number of possible explanations for this. Past experience of problems could affect access to credit if lenders use information about previous payment difficulties in their lending decisions. Alternatively, the experience of problems could lessen any stigma attached to payment difficulties and this could make the household less careful in avoiding these in the future. The evidence implies that policies addressing mortgage payment problems can have long-lasting effects.

We find no evidence for collateral effects: neither the amount of housing equity nor the presence of negative equity affects the probability of mortgage payment problems (although they will affect loss given default). This result is new and contrasts with previous work which has identified housing equity as a determinant of the aggregate level of mortgage arrears. This difference may be due to the sample period we use. The BHPS contains information on housing equity from 1993 onwards, so it does not allow us to directly measure the effects of falling house prices between 1990 and 1993 upon mortgage payment problems. It is possible that falling housing equity had already affected some mortgagors' ability to service their debts before 1993 and that these households would not appear in our sample.

We use the estimation results to construct a measure of mortgage debt at risk. Changes in the probability of payment problems and in the amount of secured debt held will both affect the amount of debt at risk. Over the sample period 1994 to 2002, we find that mean debt at risk has fallen. This implies that the probability of mortgage payment problems has fallen sufficiently to offset the effects of increasing mortgage debt over the same period. There is also evidence that mortgage debt is now concentrated in less risky households. This implies that the short-term financial stability risks associated with the stock of mortgage debt in 2002 are lower than in the mid-1990s.

Misperceptions and monetary policy in a New Keynesian model

Working Paper no. 278

Jarkko Jääskelä and Jack McKeown

Over the past decade, equity prices in the United Kingdom and other major industrial countries have risen sharply and have subsequently fallen back. Towards the end of the period in which equity prices were rising, UK household borrowing and house price inflation also picked up. One — but by no means the only — explanation for these events might be that people expected future incomes to be higher and so increased their borrowing to bring forward this higher expected future income in order to smooth consumption. But what if these expectations for future income were over-optimistic — what if the private sector expectation of higher future income were a misperception? How should monetary policy respond in a situation where behaviour today is influenced by misperceptions? In this paper, we discuss how monetary policy might react in an environment where behaviour may have been driven by over-optimistic expectations — misperceptions — about future output.

We develop a model to analyse how monetary policy might respond to these potential misperceptions about future output. Our laboratory economy is a calibrated New Keynesian model in which both the output gap and inflation depend on the expected future output gap and inflation. Both inflation and the output gap also display persistence. Misperceptions are modelled as persistent demand shocks, which feed through the expectations channel of current demand into the determination of output. It is assumed that while policy cannot create or dispel misperceptions, it can offset their effects. We assume that policy takes the form of a 'Taylor rule', that is, the central bank sets the interest rate in response to two variables — inflation and the output gap. Within this class of simple policy rules, we describe how optimal weights on output and inflation in the policy rule change in response to misperceptions. We also calibrate the costs and benefits of responding to misperceptions under uncertainty.

Using this laboratory we come to the following conclusions. First, and unsurprisingly, we find that misperceptions cause welfare to be lower. Furthermore, varying the persistence of the misperceptions, we find that welfare decreases as persistence increases — the longer agents are wrong, the worse are the effects of a given misperception. Second, by allowing for some rule-of-thumb behaviour, we look at how the degree of forward lookingness interacts with misperceptions. We find that in our set-up forward lookingness is bad — the more forward looking agents are, the more welfare is reduced by a given misperception. Intuitively, this result comes from the fact that agents who are more forward looking will try to bring forward more of their misperceived higher expected future income. Third, policy can partially offset the effects of misperceptions by responding more actively to both deviations of inflation from target and to output gaps. Because misperceptions distort demand, output and inflation are pushed in the same direction and so policy should respond more to both in order to offset the effects of misperceptions. Policy should, however, place relatively less weight on output gap fluctuations. How policy should react to misperceptions depends crucially on the persistence of the misperception — the more persistent the misperception, the less weight the policymaker should place on output. This result is intuitive; as the noise in the target variables increases the less weight should be placed on it. Fourth, we consider situations in which policymakers are uncertain about the process driving misperceptions. We find that unless the policymaker is confident there are no misperceptions, or that any misperceptions will be quickly corrected, the policymaker should assume that persistent misperceptions are present: policy rules derived on the assumption that misperceptions are persistent do better in the event that this assumption is incorrect than do policies based on assuming no misperceptions when in reality the opposite is true. When we also include policymaker uncertainty about the degree of output persistence, we note that, in the presence of misperceptions, the robust 'safe' strategy is to overestimate the degree of forward-looking expectations (and so the size of the misperception).

Monetary policy and private sector misperceptions about the natural level of output

Working Paper no. 279

Jarkko Jääskelä and Jack McKeown

There is ongoing debate about by how much the real world differs from the world described by models of rational expectations. This paper describes a simple model that offers some insight into the consequences for monetary policy design of problems the private sector and the central bank might have in estimating the natural level of output.

The paper uses a simple model with two agents, a private sector and a policymaker. The private sector bases its behaviour on its perception of the sustainable level of output and on its perception of the objectives and actions of the policymaker. The policymaker sets policy to keep inflation stable around an inflation target, and to keep output stable around its sustainable level. The paper assumes that the private sector and the policymaker have asymmetric information sets. These asymmetries cause the private sector and/or the central bank to have mistaken expectations — misperceptions — about the natural level of output. Furthermore, these misperceptions are not a function of the fundamentals contained in the model, but rather are some non-modelled factor. Three variants of the misperceptions problem are considered. In the first two cases, only the private sector has misperceived natural output, while in the third case, both the private sector and the central bank have misperceived natural output. In the first case, the private sector misperception is known by the policymaker, while in the second case, the misperception is unobserved by the central bank. In the third case, both agents' misperceptions are stochastic and unobservable (to the other agent). In all three variants it is found that, in the face of a private sector misperception, appointing a monetary policy maker who will be tougher on deviations of inflation from target than society can partially offset the negative effects of the private sector misperception.

A quality-adjusted labour input series for the United Kingdom (1975–2002)

Working Paper no. 280

Venetia Bell, Pablo Burriel-Llombart and Jerry Jones

Government policy, demographic shifts and social change have radically altered the structure of the UK labour force. For example, since the 1970s, the workforce on average has become older, better educated, and more balanced between the genders. This paper examines these changes in the labour market from 1975 to 2002, and their implications for labour quality.

Economists are interested in evaluating factor inputs (such as capital and labour) because they are measures of an economy's productive potential. The standard measure of labour is to aggregate the number of hours worked by each person in the economy. Yet this method does not take into account the fact that some people are more productive than others. By adjusting standard measures of an economy's total hours worked with a labour quality index, we can derive a truer measure of the contribution of labour to production.

An overall shift in the structure of the workforce can change the aggregate skill (quality) level in an economy. Measuring 'skill' is difficult, since it is a loose term that in part reflects the characteristics of a worker and is not directly observable. To compound matters, individuals are different, and, to a certain extent, their skill levels are subjective. For example, it could be argued that a younger workforce is likely to be more innovative and dynamic than an older one. Conversely, an older workforce, with greater work experience, might be more productive. In order to capture skill levels, it is necessary to find proxies. This paper uses information on wage differentials between worker groups as a measure of skill.

Data from the Labour Force Survey and the General Household Survey are used to construct the quality-adjusted labour input series over the period 1975–2002. The total hours worked by particular groups of workers are weighted by their respective wage bill shares. Our benchmark series takes into account gender, five age groups and four education levels.

We find that the quality of the UK workforce has increased since 1975. Adjusting for labour quality adds 0.67 percentage points per annum to the growth rate of labour input from 1975–2002. This increase can be attributed to changes in the educational distribution. Meanwhile, the workforce, in general, has become older, reflecting the temporary increase in the birth rate after the Second World War. This has had a positive effect on measured labour quality, particularly after 1990.

Adjusting labour input for quality changes has some interesting economic implications. The final section of the paper explores these issues. Its key finding is that a large proportion of what is usually considered to be TFP (total factor productivity) growth can be attributed to an improvement in the quality of the labour input. This result has no implications for the measurement of UK growth from 1975–2002, but it does help us to identify more accurately the sources of that growth.

Monetary policy and data uncertainty

Working Paper no. 281

Jarkko Jääskelä and Tony Yates

The data policymakers use to assess the state of the economy are often uncertain proxies for the things they really want to know about. Data releases referring to the most recent periods contain the most signal about the future for policymakers, but typically also contain the most noise. Many data are revised over time, and improved in the process. Those that are not revised are still uncertain, but over time other corroborative evidence arrives that can help us interpret them. We ask how policy should be designed in the face of this kind of data uncertainty. What if policymakers do not know how much variation there is in data uncertainty over time, or over vintages? We also ask whether the response of policy to what we term time variation in data uncertainty, or more properly variation across vintages, can account for the observation that interest rates seem to move more sluggishly in response to news than most models would predict they should.

We present a model that allows us to study variation in measurement error across data vintages. In our model, there are two endogenous variables the central bank has to measure: inflation, and the output gap. In the United Kingdom, and in many other countries, inflation data typically do not get revised, and therefore the measurement error in current period inflation data are (improvements in survey methods aside) the same as that in old data. Output data, however, are revised, and it is likely that early releases of output are less well measured than the revised estimates that succeed them. Our model is a metaphor for this world: inflation data are always perfectly measured, but output gap data become better measured over time.

We make three observations. First, we examine simple optimised rules for monetary policy: these rules are based on current and past-dated inflation and output data. The optimal coefficients change as the amount of noise in the output gap data increases, and as the measurement error in new data increases relative to older data. Intuitively, the more measurement error there is in the output gap data, and the worse current data are relative to lagged data, optimised simple rules put more weight on inflation compared to output gap terms; and more weight on lagged output gap terms relative to current ones.

Second, we note that an econometrician who tries to study the behaviour of central bank policy rates — but is unaware that central bank rates are designed to cope with data uncertainty — will conclude that interest rates move too sluggishly in response to news. But it is likely that vintage variation in measurement error alone cannot account for the amount of interest rate smoothing seen in the data.

Finally, we explore the effects on policy of uncertainty about noise in new data compared to that in older data. In the face of this lack of knowledge robust policies err on the side of assuming more vintage variation in measurement error, rather than less. This is an interesting theoretical result, but it could also have a practical angle: the apparent 'excess' smoothing in observed policy rates may reflect a robust response to an unknown degree of variation in measurement error across different vintages of data.

Stress tests of UK banks using a VAR approach

Working Paper no. 282

Glenn Hoggarth, Steffen Sorensen and Lea Zicchino

Stress tests were performed on the resilience of the UK banking system as part of the IMF Financial Sector Assessment Programme (FSAP). These tests revealed that the UK banking system was robust to a number of adverse shocks. Most of these tests were conducted by the large banks themselves, based on scenarios developed from the Bank of England's Medium Term Macroeconometric Model. To compare the robustness of such a conclusion to the choice of stress test, this paper proposes an alternative test of the resilience of the UK banking sector, which analyses the common developments in a measure of bank fragility and key macroeconomic variables. An advantage of the stress test proposed here is its ability to analyse — within a small system of equations — the increase in bank fragility following a shock to a single macroeconomic variable, allowing for the potential impact on other key macroeconomic variables that may also affect bank fragility. Furthermore, the test allows for feedback effects from an increase in fragility back to the macroeconomy — for example, an increase in the default rate on loans by the household and corporate sectors may cause consumption and investment to fall subsequently.

The stress tests used here, like most other methodologies, may not fully capture structural changes in the banking industry. Nonetheless, the results are robust to a number of checks and uncover some important relationships between macroeconomic dynamics and the loan write-off ratio — our measure of bank fragility. UK banks' aggregate write-offs, and particularly corporate ones, are found to be sensitive to a downturn in economic activity. Household write-offs, on the other hand, are found to be more sensitive to changes in income gearing. The results suggest that, even if the most extreme economic stress conditions witnessed over the past two decades were repeated, the UK banking sector should remain robust.

The approach to stress testing proposed in this paper is straightforward to implement and provides a useful complement to the suite of models used to assess banking sector vulnerability.

Measuring investors' risk appetite

Working Paper no. 283

Prasanna Gai and Nicholas Vause

Financial market participants and policymakers frequently cite increased 'risk appetite' as an important driver of the recent downward trend in risk premiums and yield spreads. Risk appetite reflects investors' willingness to hold risky assets and, as such, depends on their attitudes to risk as well as the size of other risks they carry on their balance sheets, such as that relating to employment. As a general determinant of asset prices, changes in risk appetite can generate correlation among the returns of otherwise unrelated assets. For example, a decline in risk appetite may help explain financial market contagion during the 1997–98 East Asian crisis. More recently, there have been concerns among policymakers that a build-up of risk appetite may have led to a 'search for yield', whereby investors demand more risky assets in pursuit of higher yields.

This paper reviews the concept of risk appetite, distinguishing it from the related notions of risk aversion and the risk premium. It suggests a precise definition of risk appetite, relating it to the price of risk — the compensation that investors require to hold a given amount of risk. This definition can be related to differences between investors' expectations about future asset returns and those implicit in options prices. Calibration of our measure of risk appetite suggests that it fluctuates within a relatively narrow range during 'normal' times, but falls sharply during crises.