Background paper for 'Economics of Credit and Debt' session at the INET/CIGI 'False Dichotomies' conference, Waterloo, Canada, 18th November 2012

Oliver Burrows[1]

Financial Stability Directorate, Bank of England

Introduction

I was delighted to receive Dirk's invitation to the session on 'Economics of Credit and Debt' at the INET conference in Waterloo. Delighted but surprised. My surprise grew more clearly into nervousness as I read Dirk's positioning paper and discovered that this was to be a meeting of scholars. It was only on the next line, where he tactfully described the group he has assembled as possessing a 'variety of intellectual pedigrees', that I began to feel that there might be room for me.

Unlike many around the table, my background is not academic. My day job in the Financial Stability directorate at the Bank of England does not afford me much time to find elegant answers to the questions Dirk has brought us together to discuss. I manage a team of analysts and spend a lot of time producing fairly short-term analysis. But what I lack in time for reflection is more than made up for in evidence from the coalface on the problems we face. I am in the fortunate position of having to think about the economics of credit and debt on a daily basis. I am constantly confronted with evidence that doesn't appear to fit into the moneyless frameworks I learnt at university. I sit in the middle of a directorate of 80 economists who come to work each morning to worry about debt and pore over the balance sheets of financial and non-financial firms. So while I fear I may not have any deep answers to the questions Dirk has laid out for us, I recognise them and very much look forward to the debate. If I can offer anything to the group, it will be to throw in some empirical evidence and some of the experience I have gained from trying to put to work some of the methodologies Dirk outlined for incorporating credit and debt into macroeconomic policy analysis.

Following Dirk's suggestion, if not his order, this paper is structured as follows: Section one briefly sets out the career path that brought about my interest in the economics of credit and debt; Section two lays out the problems I see with modern macro, particularly as judged from the perspective of my current job; Section three explains how I attempt to address this in my work; and Section four concludes.

[1] Any views expressed here are those of the author alone and in no way represent the views of other Bank of England staff.

Section one: what brought about my interest in credit and debt

My initial training as an economist – at LSE and UCL universities and then later in the monetary economics area of the Bank of England – was very much in the neoclassical paradigm Dirk outlined in his paper. It was the beginning of the noughties and despite the dramatic backdrop of the bursting of the dotcom bubble – of which I can't remember a single mention from an undergrad lecturer at university – I found myself getting used to rational agents solving optimisation problems with cute mathematical sleight of hand. Some years later I recalled going home from uni for the summer and my father asking me why he had lost half his pension and finding that there wasn't much common ground between the reasonable questions a layman might ask and the economics I was learning at uni, but I'm sure I didn't lose any sleep over it at the time.

The Bank was, and still is, a wonderful place to start life as an economist. It is a half-way house between academia and the real world; a collegiate, thoughtful environment in which to be exposed to real-world data in earnest for the first time and start to translate economic theory into policy practice. London's financial firms know this and in 2006 a friend persuaded me to leave the Bank to join the booming hedge fund industry. My work there exposed me to evidence that was really hard to explain with the toolkit I had. The Icelandic carry trade is an example that sticks in my mind: Iceland was running 20% current account deficits, single investment projects in the aluminium industry that equated to 20% of GDP were being agreed, the stock market and house prices were soaring, rates were at 15% and the currency was appreciating and I was being asked whether borrowing in euro and lending in Icelandic Krone was a safe investment. You could ignore the money and credit data and think about it all in terms of rational responses to rising energy prices (Iceland has abundant cheap power, which it exports indirectly by smelting aluminium), which means permanently higher incomes and thus higher consumption and house prices today, as well as lots of investment. But looking at a time series of the currency gave you pause for thought, as it seemed to appreciate for long periods of time in between dramatic falls. And my job was on line, or at least, it certainly felt like that at times. So I started to wonder where all that money went. If hedge funds were borrowing lending at 15% in Krone, who was borrowing at that rate and what on earth were they investing in? True, the stock market was rising at 70% a year and house prices gains were comfortably outpacing base rates, but how were they going to meet the interest payments on the debt?

Through a strange quirk of fate, the trader that ran the fund knew Wynne Godley. The connection to this mysterious academic always surprised me, as our fund was not at all academic in nature – I bought the first econometrics package when I got there and ended up only ever using it to seasonally adjust data. But they often talked on the phone and he had enormous respect for Godley's views. After unwittingly contradicting something the boss had heard on the phone and being berated for it a number of times, I

bought a copy of Godley and Lavoie's *Monetary Economics* to find out where these mysterious views came from.

Whilst my job gave me lots of questions to think about, it didn't give me much time to do the thinking. So on a holiday with a friend riding horses through the Tien Shan mountains in Kyrgyzstan, I found myself looking for a quiet spot by the fire each night to work my way through their book on monetary economics. It read like a revelation, and quite a timely one, for this was the summer of 2008. Along with many others in my industry, I had read Kindleberger's Minskian view of history in *Panics, Manias and Crashes* earlier that year and it seemed at last like I was beginning to find a home for all the awkward evidence that the past few years had thrown up.

By the end of that year, my hedge fund looked like it was teetering on the brink and I went back to the Bank, keen to get to work on fighting the crisis, but also hoping for the time to work out how Godley and Lavoie (GL) and Minsky could be put to use in practice. Four years on, the crisis seems far from over and I have not spent as much time thinking as I would have liked, but I feel like I have made some progress.

Section 2: Do we need to re-think finance in macroeconomics? One practitioner's perspective.

Dirk is pessimistic about the current state of mainstream macroeconomics. For the sake of brevity, we can crudely characterise DSGE models as the mainstream paradigm that was entrenched in academia and developed world policy institutions pre-crisis and that still largely holds sway. He points out that the early workhorse models, and some of those used by policy institutions today, have no role for credit or money at all. Exogenous shocks can be exacerbated by frictions in goods or labour markets, but absent the exogenous shocks, the models returned to equilibrium. Adding rudimentary financing to the models adds another market in which a friction could exacerbate an exogenous shock, but, as Dirk points out, this is not the same as the financial market having *caused* the instability.

Dirk is also pessimistic about DSGE as a way forwards. He proposes that 'ex-nihilio credit creation is incompatible with models which are solved by optimisation in the real sphere' and that we need a more detailed treatment of credit and debt, such that the two can't simply be netted off against each other in some sleight of hand behind a veil.

I would like to elaborate on his first point, which I think addresses the biggest problem with the mainstream approach to credit and debt (Section 2.1) and then add briefly discuss a few further points that trouble me (Section 2.2).

2.1: No role for the type of credit that account for the vast majority of lending booms and busts

As Dirk points out, DSGE models tend to either have no role for credit, or to motivate it with firms borrowing to add productive capacity to the economy. They abstract from the institutional details of the banking sector, but as I understand it, the implicit balance sheet relationship is that households lend commodity or fiat money to firms, sometimes intermediated by banks. No money is actually created by banks. In such a framework, it is possible to motivate losses on lending through the failure of a fraction of the borrowers' projects – i.e. simple credit risk. Without explicitly modelling risk, some models find novel approaches to introducing this sort of risk, such as a recent popular

attempt at introducing a banking sector into a DSGE setting in which banks 'steal' deposits from depositors each period with some fixed probability (i.e. Gertler & Karadi, 2011). This is surely useful progress for the DSGE literature, but I think it is inadequate in two key respects.

Borrowing for the purchase of existing assets must be modelled

First, loans to firms undertaking investments that intend to expand the supply capacity of the economy and thus can be expected to yield a return capable of repaying the debt ('credit for real investment' hereafter) accounts for a surprisingly small proportion of borrowing. For the UK, looking only at bank loans, I estimate that it accounts for at most 15% of the existing loan stock.[2] The rest is lending to firms and households for the purchase of existing assets (primarily houses 65% and commercial property 13%, hereafter 'credit for asset purchases') and for consumption 7%.

Furthermore, I strongly suspect that it is credit for asset purchases that accounts for the lion's share of financial instability. From a theoretical perspective, Minsky's *Financial Instability Hypothesis* (1992) has a particularly strong appeal for lending for asset purchases due to the immediate positive feedback between lending and asset prices. If agents borrow to buy an asset expecting to make capital gains and the act of buying the asset pushes up the price, it can quickly become self-reinforcing. By contrast, the feedback in borrowing for real investment is likely to by much weaker, if present at all, as the return to the projects is far slower and needs to be matched by demand for the good or service produced, rather than simply a rising asset price.

From an empirical perspective, taking the credit boom in the UK between 1997 and 2008, the vast majority of the rise in bank lending supported the purchase of assets (at least 78%), with 'credit for real investment' contributing 8% and barely outpacing income growth over the period. And looking at the losses suffered by the UK banks since the crisis, commercial property stands out on the weaker banks' loan books as not only the largest source of losses, but also the worst losses relative to the capital held against the lending pre-crisis – i.e. the type of lending where they most misjudged the underlying risk.

Banks' role in creating credit inside money, rather than intermediating outside money, must be modelled

Second, by ignoring banks' ability to create credit (and instead positing that they simply intermediate existing outside money), they ignore the crucial role banks play in setting monetary conditions and thus an important structural role banks play in the economy through their ability to impact asset prices and interest rates and in turn the cross-sectoral allocation of resources and intertemporal spending decisions. Recent work by Cecchetti (2012) paints an interesting picture of credit growth reducing productivity growth by misallocating labour and capital towards the financial system.

[2] For this calculation, I just consider the stock of UK-resident bank lending to households and private non-financial firms (PNFCs) and attribute lending to be either 'credit for real investment' or 'credit for asset purchases' based on the activity classification of the borrower. Mortgage lending to households accounts for 65%, unsecured lending to households 7%, lending to firms that own and rent out but do not develop real estate 13% and lending to other non-financial firms, which I classify as 'credit for real investment', 15%. Note that this is likely an overstatement. As well as credit-driven bubbles in housing and CRE in the UK in the pre-crisis years, there was a bubble in the private equity market, in which private equity firms bought PNFCs at ever higher debt/income multiples. The debt for the purchases was taken on by the target company – not the PE firm – and thus shows up in the non-CRE PNFC sector. Note also that adding PNFC bond issuance increases the figure to 21%, while adding lending to financial firms is trickier, but on a crude measure brings it down to 11%.

And the work of Mian and Sufi (i.e. 2009) has established a clear role for easing credit conditions and a subsequent rise in leverage pre-crisis in explaining the weakness of consumption in the United States during the crisis.

This may have important consequences for monetary policy and, excitingly, may provide a useful theoretical underpinning for the introduction of 'macroprudential policy'. A broad consensus is emerging in the developed world around the need for macroprudential policy as a response to perceived flaws in the pre-crisis framework of inflation targeting and institution-focused microprudential regulation.[3] There is a case to be made that this lack of attention stems directly from a flawed approach to monetary economics. In the debate around macroprudential policy, much has been made of the lack of attention to the externalities to bank lending and the systemic risk to which they can lead. And there is a perception that macroprudential policymakers face a trade-off between resilience and growth in the short-run – increasing capital requirements increases resilience, while decreasing it increases lending growth which in turn increases supply growth.[4] But to evaluate this perceived trade-off, one needs a model with a structural role for credit in the economy that captures both the threat credit poses to resilience – and 'credit for asset purchases' is likely crucial here – and the beneficial role credit plays in facilitating growth.

2.2: Some other problems to consider

Equilibrium

The very concept of an economy always tending towards equilibrium seems problematic to me. As Dirk points out, it leaves crashes to be explained as exogenously triggered events, exacerbated by frictions. This leaves people in my line of work – financial stability policy – without a model for financial crises; not only do we not know what causes them, but we don't understand how they work, which makes policy prescription difficult.

More worryingly, I think equilibrium models can have a dangerous impact on the way we interpret what is going on around us. If the framework in which you think tends to equilibrium, then you are likely to look for reasons that the outcomes you observe in the real world are points on a stable, sustainable path. As evidence of this, I'd cite the widespread use of a new era of monetary stability and low real rates to explain the rapid rise in house prices and debt seen in much of the developed world before the financial crisis.

Representative agents

I have found stock-flow consistent models a very useful framework for understanding the limitations of representative agent. A household that owns a bank and simultaneously borrows from and deposits money in that bank is not going to worry about the bank's solvency. The household's debt may rise to thousands of times its income, but the interest payments it makes on the debt are clearly identically offset by the interest it receives on its deposits and any margin the bank makes, which is paid back to the household in dividends. Debt levels are only anchored in the real world by a fear of default. Modelling this necessitates a separation of borrowers and lenders.

[3] See, for example, Galati and Moessner (2011).
[4] See, for example, Jacome and Neir.

The wrong sort of micro-foundations? Contrasting DSGE to ABM

DSGE models are often described as 'micro-founded', in the sense that the behaviour of sectors is representative of the agents it comprises, i.e. the economy is governed by the solution of optimisation problems that are set out as if they are faced by individual households and firms. It is interesting to contrast this to the sort of micro-foundations used in agent-based modelling by other branches of social and physical science, in which there is an explicit attempt to model the behaviour of individual agents and the manner in which they interact, before the behaviour of the system is simulated and the macro outcomes are observed. If the agents interact in an interesting (generally non-linear) fashion, then the macro outcomes cannot be easily calculated analytically from the specification of the micro behaviour. It strikes me that this places a strict criteria for success on the modeller: the model must be able to describe the micro behaviour *and* the macro behaviour. Judged against this criteria, DSGE fails – households and firms clearly don't solve optimisation problems over infinite horizons.

My area of work, financial stability analysis, necessarily involves both micro and macro perspectives. Macro dynamics in credit, asset prices and income drive booms and busts, but when the fragility crystallises, it does so on the balance sheets of the weakest banks and firms. Firms that differed in their micro characteristics from their competitors. Furthermore, once a crisis starts, there is enormous interest in the mechanisms through which it is propagated. This manifested itself in policy circles post-crisis in modelling the financial system as an interconnected network.[5] Policymakers are already coming to accept that the network approach offers explanations for the propagation of shocks and thus policy prescriptions that are not available in mainstream macro models.

Agent-based modelling offers the possibility of going beyond modelling propagation of shocks across static and arbitrarily generated networks and model the underlying complex systems from which they arise. As Dirk points out, it holds the prospect of uniting micro and macro and explaining financial crises as the sort of phase transition seen in other physical and social systems. In time, it might lead to a different understanding of crises, which in turn would require different models for evaluating policy.

Section 3: How I try to address these challenges

My work on credit and debt has taken a number of different paths. Much of it has focused on analysing UK balance sheet data in a fairly informal manner, looking for Minsky-style stories. This is described in my 2011 paper with Richard Barwell, which focuses on sectoral flow-of-funds (FoF) data from the national accounts. Since then, I have tried to expand the data work by: (i) broadening out the flow-of-funds data to capture the financial sector in more detail; (ii) disaggregating the sectoral data down to an individual firm level to better identify where and how financial fragility crystallises. On the modelling side, I have focused on building SFC models based on the UK sectoral balance sheets.

[5] See, for example, Haldane (2009) and Gai and Kapadia (2010).

Section 3.1 discusses the data work; Section 3.2 turns to the SFC modelling approach and its extensions; and Section 3.3 offers some thoughts on where this does and doesn't meet the criticisms outlined in Section 2.

Section 3.1: A SFC approach to the UK data

In my 2011 paper with Richard Barwell, we draw together the FoF data from the UK national accounts into a slightly simplified accounting framework and re-examine the build-up to the 2008 financial crisis in light of the balance sheet developments of the time.[6] The paper starts with the observation that despite the 15 years preceding the crisis being some of the most stable in the UK's history in terms of output growth and inflation – earning it the moniker 'The Great Moderation' – balance sheets and asset prices were anything but stable over this period (Charts 1 & 2). It goes on to argue that the balance sheet developments of the day, primarily an enormous expansion of bank credit and inside money, had little obvious impact on contemporaneous output and inflation and thus did not attract sufficient attention from economists who viewed the world through the prism of mainstream macroeconomics. We further argue that the interpretation of the Dotcom bubble and bust – that the ensuing recession was small and it was possible to 'mop up after' asset price bubbles burst with accommodative monetary policy – was misleading precisely because it failed to understand the difference between credit- and equity-driven asset price bubbles. Looking instead from a flow-of-funds perspective, we claim that it would not only have been harder to ignore developments in balance sheets and asset prices, but also that it would have been easier to link them to some of the interesting puzzles of the day.

Chart 1: UK output and inflation

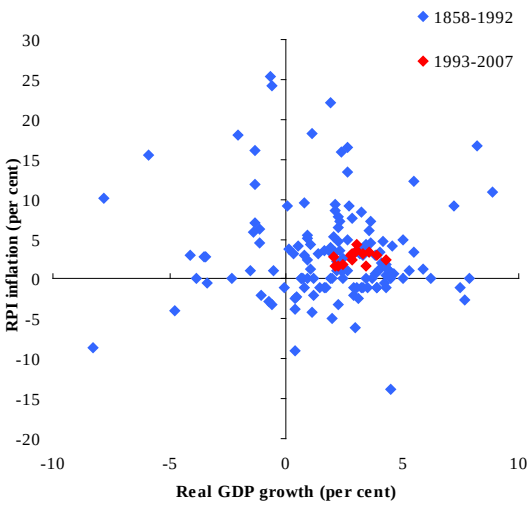

Source: Bank calculations

Chart 2: UK sectoral debt & asset prices

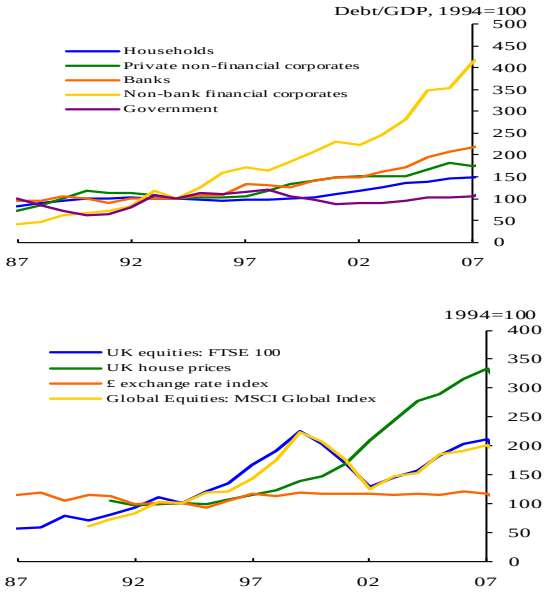

Source: ONS, Datastream and Bank calculations

[6] For any readers unfamiliar with the flow-of-funds accounting framework, please see Appendix 1.

The paper attempts to explain a lot of developments of the time, but I will focus here on developments in the household sector (Section 3.1.1), before turning briefly to the corporate sector (3.1.2).

3.1.1: The household sector

After a period of relative calm in the wake of the early 90s recession, household balance sheets grew in line with income in the mid 90s (Chart 3). But from around 1997 onwards, they started to grow much more rapidly, as households borrowed from banks. The counterpart financial assets were deposits and claims on pension funds (Chart 3). But revaluation effects were having a far larger effect on household balance sheets as house prices, and to a lesser extent equity prices, rose (Chart 4).

Chart 3: Household net lending and net accumulation of financial assets and liabilities

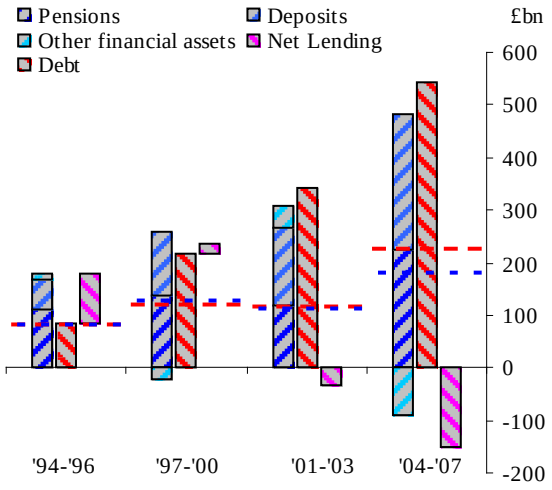

Source: ONS and Bank calculations

Chart 4: Household balance sheet growth

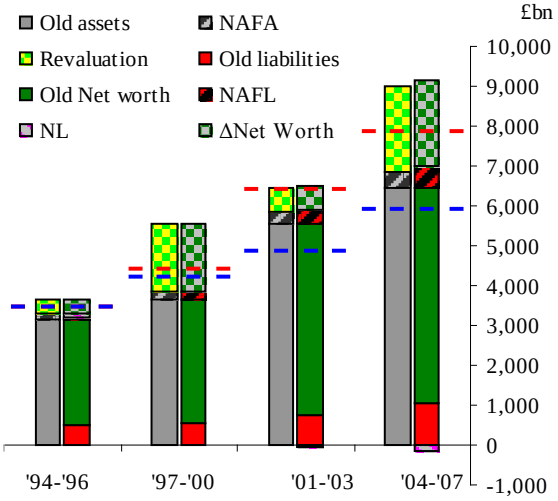

Source: ONS and Bank calculations

Dashed lines give benchmarks for the growth in stocks that would occur if they grew in line with trend nominal GDP (5%) in each period considered (red line), and cumulatively from 1994 (blue line). Solid bars represent stocks, dashed bars show flows and hatched bars show revaluation effects and changes in net worth.

Developments in house prices and household debt attracted plenty of attention over this period. The view generally taken was that a substantial part of the rise in house prices was to be expected given the macro developments of the day: long-term real interest rates had fallen over the course of the nineties; the availability of credit had eased significantly; inflation had fallen to low and stable levels not seen in a generation, significantly reducing households' income gearing in the early years of their mortgages; and the rate of household formation appeared to have outstripped the rate of housing supply for some time, for a variety of structural reasons.[7] Because various attempts to quantify the impact of these factors suggested they may fall short of fully explaining the rise in prices economists remained open to the possibility that the rise in house prices contained a bubble-like element.[8] Given the rise in house prices, the rise in secured

[7] See, for example, the May 2004 MPC Minutes, Bean (2004), Nickell (2004).
[8] See Weeken (2004) for an attempt to quantify the impact of lower real rates on house prices.

debt was easily explained.[9] Most attention focused on the impact of house prices on consumption. Although some causal channels were noted, particularly the availability of collateral against which to borrow, the impact of higher house prices on consumption was generally argued to be muted, largely because the positive wealth effect on households that owned sufficient housing for their lifetime needs was offset by the negative wealth effect on those that did not. Lower real rates and some exogenous structural factors, like credit availability and the rate of household formation were thought to have led to a one-off increase in house prices, but the impact on macro flows, such as consumption, was expected by many to be muted.

A more puzzling development, in the United Kingdom and elsewhere, was the growth of global trade imbalances. Most developed Western countries ran pronounced current account deficits with most developing Asian economies and commodity exporters ran counterpart surpluses Global interest rates fell to historically low levels. This development was awkward for mainstream macro models, which for the most part predicted that the rapid opening up of Asian markets and attendant increase in labour supply should increase global interest rates and cause Asia to run trade deficits with the West, as capital poured from West to East to fuel the investment boom in the East.[10] Quite the opposite seemed to be happening. A variety of explanations arose to explain the apparent anomaly, with the hypothesis of a 'savings glut' in Asia and a lack of credible financial assets in Asia amongst the most popular. Policy makers internationally, ascribed some weight to these views, and often voiced concern about the possibility of a fall in asset prices should the global imbalances unwind in a disorderly manner.[11]

Linking together the stocks and flows

A balance sheet accounting framework offers an interesting perspective on the period. Over the period 2001-7, the household sector ran a net lending deficit of £175bn, which roughly coincided with a RoW net lending surplus (i.e. a UK trade deficit) of £186bn. But households did not borrow directly from foreigners; instead, they borrowed from banks. And their borrowing from banks, £782bn vastly exceeded their cumulated net lending deficit (i.e. it banks were creating credit over and above intermediating credit between sectors). Households' borrowing from banks also exceeded their increased deposits with banks, causing banks to run a 'Customer Funding Gap' (CFG) of £412bn, i.e. a growing gap between their deposits from customers and their loans to customers. This was a break with historic behaviour and was met by banks by a large expansion of wholesale funding, particularly securitisation. A substantial proportion of this funding appears to have been supplied by foreigners: over the period, RoW is estimated to have purchased £647bn of UK private sector debt, while the bank and OFI sector issued £866bn in combination and the non-financial corporate (NFC) sector only £111bn. The bank and insurance company and pension fund (ICPF) sectors appear to have bought the rest of the issuance. The increasing reliance of banks on particular wholesale funding markets was noted with growing concern in the Bank of England's *Financial Stability Reports* of the time.

[10] See, for example, Spange (2007) for a view on the theory and evidence.
[11] See, for example, King (2006) and Bank of England *Inflation Report*, February 2006; Trichet, 4th May 2006 ECB press conference.

Diagram 1: the impact of household behaviour on balance sheet growth

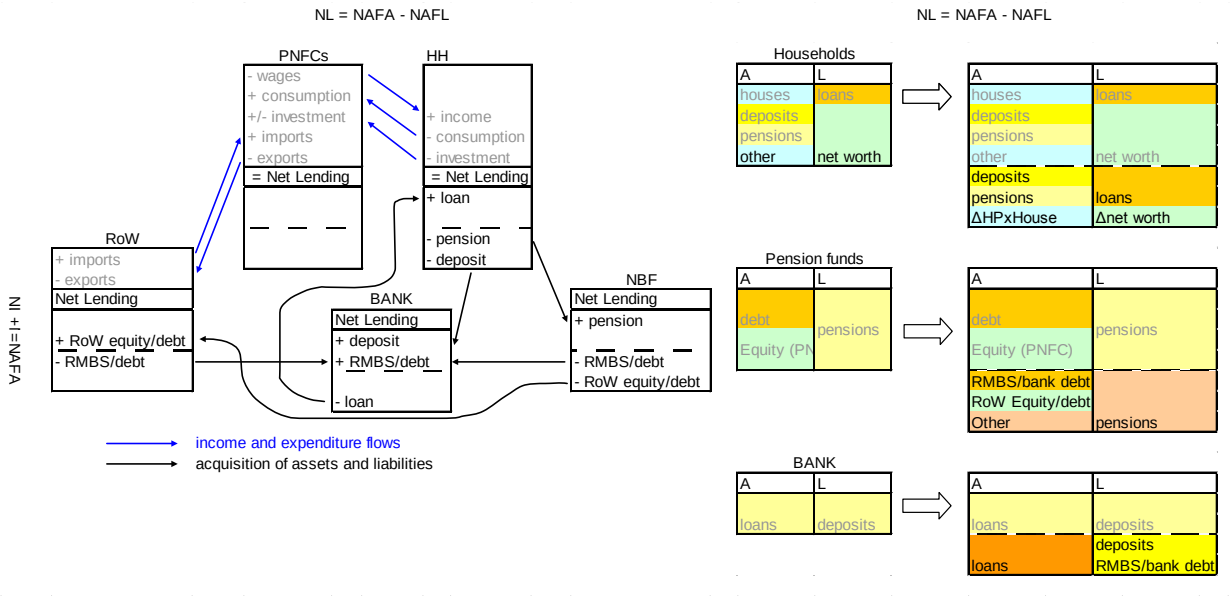

In this example, the HH sector runs a net lending deficit by spending more on consumption and investment than it earns. This causes it to run a deficit with the NFC sector, which in turn runs a deficit with the RoW by importing more than it exports – assuming flat government and financial sector net lending balances, this leads to a trade deficit at a national level. To fund its financial deficit, the HH sector's NAFL must exceed its NAFA; thus its loans from the banking sector exceed its purchases of pension assets and deposits. The excess of loans over deposits leads to a CFG for two reasons: first, because the net lending deficit requires that NAFL exceeds NAFA; and second, because the HH sector splits its NAFA between deposits and other financial assets – in this case pension assets. In the first case, the funds have flowed from the banking sector to the RoW and will very likely return as wholesale funding (unless the RoW buys an asset from the HH or PNFC sector and that sector deposits the proceeds with the banking sector) – in the example, they return directly as a purchase of bank RMBS/debt. In the second case, the funds have flowed to the NBF sector and will also likely return as wholesale funds – in the example, they again return directly as RMBS/debt.

While a more formal model is required to identify what was going, a story about a housing bubble seems to fit these stylised facts. When young households borrow from banks to buy housing from old households, old households receive funds to consume or invest (Diagram 1). They saved most of the funds, some in the form of bank deposits, some via claims on ICPFs and some in other financial assets. Note that at an aggregate level, the household balance sheet expands: liabilities increase by the amount of the loan, and assets increase by the same amount, comprising the financial assets that the older household chooses to buy (deposits, pension assets, mutual funds shares, etc.). If older households choose to hold some of the funds in assets other than deposits, then fewer funds are deposited with banks in the form of deposits than are lent out as loans, causing a gap between loans and deposits to arise. The banking sector finances this through increasing its wholesale liabilities (interbank loans from foreign banks, bonds, securitisations, etc). Between 2001 and 2007, the issuance of Residential Mortgage Backed Securities (RMBS) played a large role in closing banks' funding gap, with the bonds sold to domestic and foreign bank and non-bank financial institutions. Direct funding from foreign banks also appears to have been important.

It seems quite plausible that at some point a feedback loop formed between borrowing and house prices. As households anticipated house price gains, they were willing to borrow more and pay more for housing, further bidding up house prices. As house prices rose, the loan-to-value ratios on existing lending declined, making banks increasingly willing to lend, particularly against housing collateral. Meanwhile, the saving rate fell, perhaps in response to a perception of higher real wealth, and the current account deficit widened.[12] The flow of funds from foreigners to banks, or securitisation vehicles, provided more than

[12] See Davey (2001) for evidence that falls in the savings ratio in the late 90s and early 2000s were driven by rising wealth.

sufficient financing for this deficit, with the excess being part of the large cross-border flows of capital between non-bank financials, corporates and governments in the United Kingdom and abroad.

This story receives some support from disaggregated data. Chart 5 reveals the marked cross-cohort variation in the change in the size and structure of balance sheets across this period. Young households' net financial wealth fell as they took on more debt to buy housing, and older households' net financial wealth rose as they sold housing to younger generations in order to buy financial assets. Housing wealth rises most for middle-aged households, who hold most of the stock of housing.

And there is some evidence from aggregate data that at least some households expected house prices to keep rising. The profit an investor expects to make from a buy-to-let (BtL) investment, where a household buys a property for the purpose of letting it out to tenants, is roughly equal to the expected price appreciation, plus the rental yield less the mortgage cost.[13] Chart 6 shows aggregate data on rental yields and mortgage costs.[14] It suggests that as house prices rose, BtL investors became increasingly willing to rely on expected increases in house prices to make their investment profitable.

Chart 5: Buy-to-let rental yield

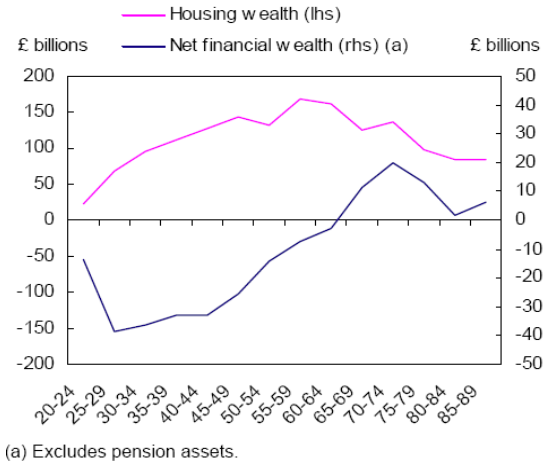

(a) Excludes pension assets.

Source: British Household Panel Survey and Bank calculations

Chart 6: Loan to income ratios

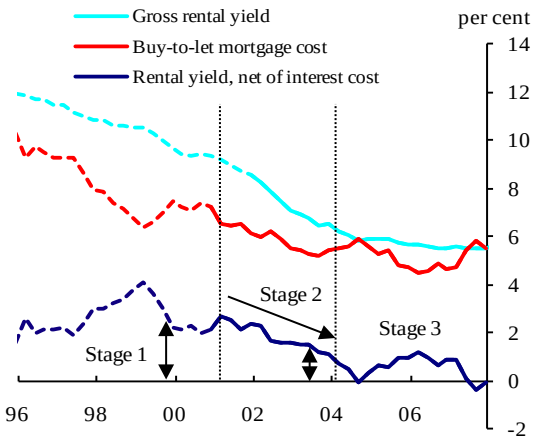

Source: Bank of England, ONS, Bank calculation

For highly leveraged BtL investors, this appears to be reminiscent of the three stages of the financing of a bubble in Minsky's Financial Instability Hypothesis. In the initial stages of a credit expansion, investors borrow only what they can repay with their expected income from the project (which seems plausible in the late 90s). In the second stage, as the expansion rolls on and expectations become more bullish and lending standards relax, agents borrow as much as they expect they can service with their expected income (allowing for some operating costs, this seems to be roughly the case for the early noughties). In the final stage, borrowers borrow more than they expect to be able to service with future income, relying instead on capital gains to cover their borrowing (this is true by the end of the period, when rental yields are negative, and probably true of much of the second half of the noughties, if operating costs are non-negligible). The

[13] In reality, investors must also allow for operating costs, primarily the risk that the house may lie empty for some periods. Taxes also play an important role, as interest cost is tax-deductible. This means that any deficit of rent below interest cost can be offset against profit elsewhere to reduce the investor's tax burden, complicating the calculation.

[14] Rental yield data are only available back to 2002, and mortgage cost data to 2001. Prior to that, RPI rents data and house prices are used to calculate a rental yield, and a spread of 100bp over the observed 95% LTV, 2yr mortgage rate is used for mortgage cost. For all periods, the mortgage cost is calculated as a spread over 2yr swaps and then added to a 20yr swap rate, to give a better measure of expected interest cost over the life of the investment.

same principle of over-borrowing and relying on price appreciation appears to have been true of some elements of the owner-occupier market, with FSA data suggesting that repossessions have been concentrated in households with very high initial loan-to-value ratios and self-certified mortgages.

Household balance sheets and financial fragility

Viewed as unconnected phenomena, the rise of house prices and household debt, the current account deficit and the growing reliance of banks on short-term wholesale markets and foreign funding each seemed like small risks. Importantly, while each seemed unable to go on forever, they also seemed unlikely to come to an abrupt halt of their own cause, in what appeared to be a very benign macro environment. But viewed as interconnected phenomena, and with the benefit of hindsight, it is clearer that their eventual correction was inevitable. Household debt cannot grow faster than income forever – households simply run out of income to service the debt. While it may have taken many years for households' appetite for housing to have declined of its own volition, a steady supply of lending from the banking sector was required for the level of house prices *not to fall*. And for the banking sector to supply that lending, they in turn needed a steady supply of funding from the wholesale markets, and particularly foreign investors in those markets, upon which they had become increasingly reliant.[15] The nature of banks' balance sheets – an increasingly leveraged collection of loans, financed with significant amounts of wholesale funding – meant that the supply of credit could be contracted very quickly if wholesale investors became worried about households' ability to service their debt and thus about the solvency of banks.

Viewed with the benefit of hindsight as interconnected phenomena, the rapid expansion of household debt, with its counterpart in increasingly stretched bank balance sheets, looks much like many of the great credit bubbles of the past. As Kindleberger concludes in 'Manias, Panics and Crashes', a pattern of increased investor optimism, declining risk-aversion on the part of lenders and the resultant wave of leveraged investment for short-term capital gains rather than for the returns associated with the productivity of the asset ran through many of the crises of the past 400 years. And while some of the shocks that brought an end to these booms were unpredictable, others were highly probable: "At some stage it was inevitable that the lenders would reduce the rate of growth of their loans to these increasingly indebted borrowers, although the details and the timing of these moves could not have been predicted."

3.1.2: The corporate sector

The paper delves into the private non-financial corporate (PNFC) sector in some detail. To summarise very briefly, PNFC balance sheets grew rapidly in the late 90s during the dotcom bubble (Chart 7); growth slowed as corporates retrenched in the global slowdown of the early 2000s; and then they grew quite rapidly in the period running up to the crisis. This section addresses the two episodes of expansion in turn.

[15] See Speight and Parkinson (2003) for a discussion of bank wholesale funding activity to support household lending.

Chart 7: Corporate NAFL/GDP

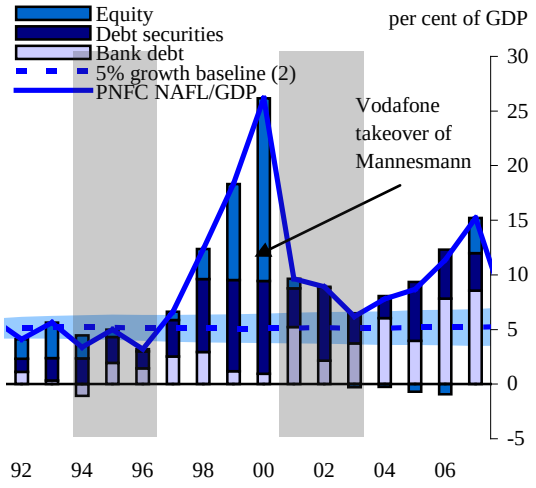

Source: ONS and Bank calculations

Chart 8: Buyout deals

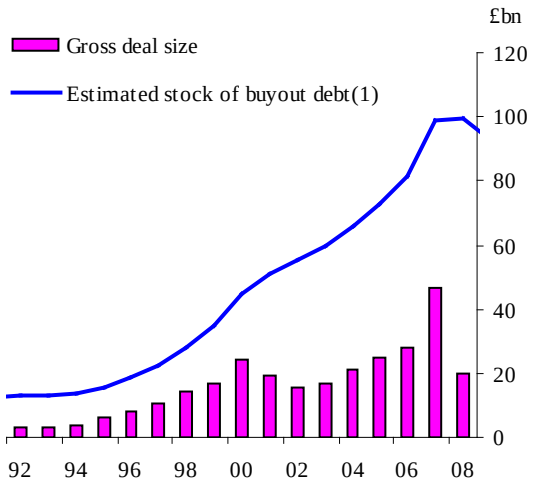

Source: ONS, Bank caculations

Late 90s Dotcom expansion

Most of the balance sheet expansion in the UK in the late 90s was driven by mergers and acquisitions, particularly in the telecommunications industry. Ever higher prices were paid for telecoms companies, with the £100bn acquisition of Mannesmann by Vodafone in 2000 epitomising the balance sheet expansion of the day. Crucially, with equity markets buoyant, most of the expansion was financed with equity, although there was a strong pick-up in corporate bond issuance. Bank credit played little role. This was important in the collapse of the dotcom bubble, as there appeared to be little impact on demand in the UK. Bank credit supply was little affected and any wealth effects from falling pension values or direct holdings of equity and debt were muted.[16] Perhaps the most important impact of the dotcom crash was the confidence it gave policymakers in dealing with asset price bubbles.

Policymakers debated the possibility that conditions in financial markets could spill over into the real economy, but opinion was divided over the optimal policy response. The prevailing wisdom, espoused by Alan Greenspan among many others, was that central banks should not use monetary policy to intervene pre-emptively – to prick a bubble before it poses too big a problem; better to let events run their course, and 'mop up' if and when any bubble burst. Blinder and Reis noted in 2005 that:

> 'This "mop up after" strategy received a severe real-world stress test in 2000-2001, when the biggest bubble in history imploded, vaporizing some $8 trillion in wealth in the process. It is noteworthy but insufficiently noted, that the ensuing recession was tiny and that not a single sizable bank failed. In fact, and even more amazingly, not a single sizable stock brokerage or investment bank failed, either. Thus the fears that the "mop up after" strategy might be overwhelmed by the speed and magnitude of the bursting bubble proved to be unfounded ... If the mopping up strategy worked this well after the mega-bubble burst in 2000, shouldn't we assume that it will also work well after other, presumably smaller, bubbles burst in the future?'

[16] For more evidence on why wealth effects might be muted, see Starr-McCluer (1998).

This willingness to leave the market to its own devices had a parallel in the regulatory world, where the Turner Review described the UK regulator's approach, before the crisis, as:

'...based on a sometimes implicit but at times quite overt philosophy which believed that:

- Markets are in general self correcting, with market discipline a more effective tool than regulation or supervisory oversight through which to ensure that firms' strategies are sound and risks contained.'
- The primary responsibility for managing risks lies with the senior management and boards of the individual firms, who are better placed to assess business model risk than bank regulators, and who can be relied on to make appropriate decisions about the balance between risk and return, provided appropriate systems, procedures and skilled people are in place.

2003-2007 credit expansion

Coming out of the 2001-3 slowdown, investment growth did not pick up very strongly and PNFCs were measured as running a historically unusual net lending surplus in the pre-crisis expansion.[17] Nonetheless, their aggregate balance sheet was expanding considerably faster than output growth, this time driven by bank lending (Chart 7). In turn, this was driven by two key stories: rapidly expanding credit and valuations in the commercial real estate (CRE) and private equity (PE) markets (Charts 7 & 8).

In the CRE sector, the story is similar to that illustrated for the buy-to-let segment of the residential market. CRE firms borrowed money to buy property, predominantly in secondary markets rather than through development, and became engaged in a process of bidding up prices, which again led to positive feedback: rising prices increased their net wealth and eased their borrowing constraints. Credit conditions eased, with LTV ratios and interest-cover ratios falling from 2003-2006, partly driven by financial innovation opening up the market to a broader spectrum of investors.[18] As with residential property, there was a significant impact on balance sheets, as CRE firms took on newly created bank debt to buy property and the funds found their way back to banks' balance sheets, likely in part as wholesale funding. But there was no obvious impact on activity or consumer prices: wealthy individuals and ICPFs enjoyed an increase in wealth as property prices rose, but any impact on consumption was never likely to be noticeable.

Private equity-sponsored buyouts again had counterparts in the balance sheets of banks, high-wealth individuals and institutional investors like pension funds. High-wealth households and pension funds invested in private equity firms, which used the funds as an equity stake in a leveraged purchase of the equity of a PNFC, with the debt provided by banks – and generally sold on to other banks through syndication. Because the existing equity of the target firm was being purchased with a mixture of equity and debt, the target firm ended up more leveraged. The transactions brought into sharp relief the option-like payoff of equity investments: if the firm proved profitable, the private equity investors earned dividends on their investment and were compensated for their risk or floated the firm at a profit; if it did not, the equity stake was wiped out and the company was turned over to the debt-holders.

While buyouts received lots of press coverage, balance sheet restructuring also occurred through equity buybacks, which became popular at the time. This may have been driven by the availability of cheap debt

[17] Although interestingly, this could well be due to mismeasurement. See August 2012 *Inflation Report*.
[18] See, for example, the Bank of England's December 2005 *Financial Stability Report*.

financing leading to firms choosing higher leverage voluntarily, or it may in some cases have been to ward off potential purchasers, who would otherwise have been attracted to leverage it up and extract cash. PNFC net equity issuance actually turned negative over the period (Chart 7). It is not clear from the data whether the buybacks were financed out of savings or debt, but either would have contributed to an increase in leverage, at least at book cost. There is again little reason to expect an impact on macroeconomic flows from LBOs or share buy-backs, as the beneficiaries were again high-wealth households and pension funds – neither of which have strong channels to aggregate consumption. The funds likely ended up in other financial assets, once again pushing up asset prices and possibly contributing to a rising funding gap at banks.

There is thus a common theme running through the main stories behind corporate sector balance sheet expansions – a period of corporate debt expansion outpacing income, but with the debt financing the acquisition of commercial property (inflating property values in the process) or financial assets (equity). While this appears to have had little effect on macroeconomic flows, it had a significant effect on balance sheets. Not only did aggregate balance sheets grow across sectors, but the distribution of assets and liabilities within those balance sheets, particularly the corporate and banking sectors, made for a much more fragile system. A tail of highly indebted corporates arose, particularly in the real estate and LBO sectors. Their fragility was a credit risk to the banks, who financed their leveraging and became highly exposed to the value of their assets and income streams. While the equity investors in both types of transaction tended to enjoy increases in net wealth, at least in the short-term, this wealth was not available to support future financing problems of the indebted corporates.

3.2: Simple stock-flow consistent models of the UK sectoral balance sheets

Given the evidence set out in our paper, I set out to build models that could explain some of the balance sheet expansion and associated asset price rises and consequent crashes. I was initially interested in trying to explain the dynamics of the housing and commercial property markets and the response of the corporate sector to the availability of cheap bank lending, via LBOs and simply by choosing to increase gearing. While I made some progress on this in really simple, stylised models, I dropped this work to try to build a baseline SFC model for the UK which we could use for macroeconomic scenario analysis. The aim was to ground it in UK sectoral balance sheets and cover them in sufficient detail to be able to answer questions like: "what would happen if foreigners' appetite for UK bank (or govt) debt fell sharply?" or "what would be the impact on sectoral balance sheets of a change in regulation that forced ICPFs to hold more debt?".

The models have been based on Godley and Lavoie (2007).[19] I do not have any finished models to date, but can present work in progress on the benchmark model, which may be of interest to others as I am not aware of anyone having tried to build this sort of SFC model using real-world data. Section 3.2.1 introduces the UK FoF data and explains the choices I make to reduce the dimensions of the UK national accounts for the purpose of the model I am currently building. Section 3.2.2 briefly explains the sort of behavioural

[19] For any readers unfamiliar with the flow-of-funds accounting framework, please see Appendix 1.

equations I have been working on and my attempts to estimate and calibrate bits of the model. Section 3.2.3 explains further work I need to do to complete my current model.

3.2.1 The UK flow-of-funds data

As a legacy of a time of greater interest in the flow of funds, the ONS publish all of the data I need to calculate a closed system of accounts in their Quarterly *National Accounts* publication. Sectoral accounts are available for the seven high-level sectors described in the national accounts: Households, Non-financial corporates (NFCs), Monetary Financial Institutions (MFIs), Insurance Companies and Pension Funds (ICPFs), Other Financial Institutions (OFIs), Government and the Rest of the World (RoW).[20] Output, expenditure and income flows are given in far greater detail than is necessary for this model. The sectoral *financial accounts* and *balance sheets*, which together comprise the UK equivalent of the dataset published separately in the US as the *Flow of Funds*, are also described in more detail than is necessary for this model. Assets and liabilities of the sectors are broken down into 31 line items. While not all 31 x 7 entries in the matrix contain data, it is nonetheless too rich a description of the financial system to be easily modelled, so I aggregate across asset classes and sectors to reduce the dimensions of the dataset.

Aggregation of the balance sheet data

Much of the detail of the flow-of-funds data lies in breaking down broad asset classes (such as 'securities other than shares') into more detailed sub-classes (such as 'money market instruments' and 'bonds') and then further breaking them down into the sector for whom they are a liability (such as 'bonds issued by UK local authorities'). We can reduce much of this detail by narrowing the set of assets to six simple classes – deposits, loans, bonds, equity, pensions and physical assets – where 'pensions' is used as a generic liability of the ICPF sector, which can be thought of as a long-term investment product sold to the household sector for savings purposes.[21] I further simplify the balance sheets of each sector to preclude them from holding certain instruments, as detailed in Table 1. Most of the simplification is not economically meaningful – for example, I do not allow PNFCs to hold bonds, but this is of little consequence as their bond holdings are likely to be very small. A more meaningful restriction is to exclude households from directly holding equity or debt, with all their non-deposit financial assets managed by the ICPF sector. While unrealistic for countries like the US, this is not too strong an assumption for the UK, where bond and quoted equity holdings were estimated at £207bn in 2010, compared to claims on ICPFs of £2,270bn and total financial assets of £4,347bn.

[20] See the ONS' *Quarterly National Accounts* and the annual *Blue Book* for the raw data.

[21] It seems reasonable for the purposes of this model to think of a single, long-term liability to cover pension funds and life insurance. This classification implicitly ignores insurers' general insurance business. At end 2010, ONS data suggests general insurance funds managed assets totalling £107bn, compared to £1,324bn for life insurers and £1,289bn for pension funds, so this omission is unlikely to be too important.

Chart 9 shows the balance sheets of the seven core ONS sectors that arise from my choice of asset classification. It is clear that there are large residuals for the MFI, OFI and RoW sectors. Chart 10 attempts to break down these residuals. It is clear that the residuals mainly comprise claims of the MFI, OFI and RoW sectors on each other and are largely intra-financial sector claims. These claims arise in part because of the unusually high concentration of global financial activity in London, relative to the size of the UK economy, with the largest share accounted for by derivatives activity. They are further exacerbated by the national accounts being constructed on a residency, rather than an ownership basis: this means that UK-resident, foreign-owned banks' subsidiaries balance sheets are picked up in the MFI sector and intra-group activity is captured as claims flowing between MFIs, OFIs and RoW. I suspect this problem applies to many other countries too.

Chart 9: UK sector balance sheets

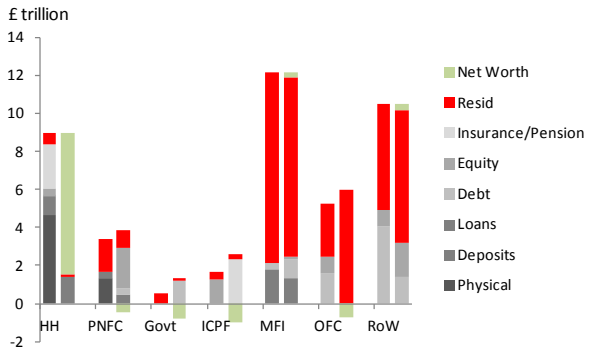

Source: ONS and Bank calculations

Chart 10: MFI, OFI and RoW balance sheets

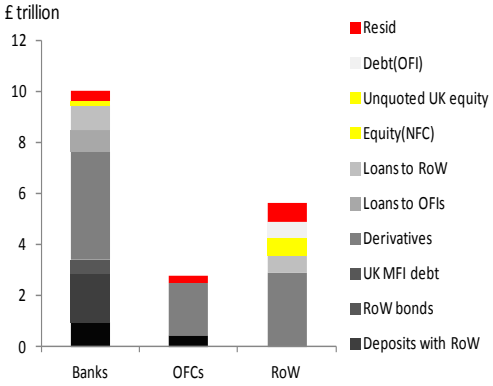

Source: ONS and Bank calculations

For the purposes of this model, which focuses on the use of financial markets by the real economy, I think that it is useful to abstract from much of this intra-financial sector activity. To achieve this, I hypothecate the MFI sector into a 'domestic banking sector' (BANK) and an investment banking sector and add the investment banking sector to OFI and RoW to construct a 'hot money sector' (HOT). The domestic banking is constructed to holds the loans of the MFI sector to households and NFCs and a corresponding proportion of MFIs liquid assets (in the form of gilts). It takes an appropriate proportion of the MFI sector's equity and all of the MFI sector's deposits from households and NFCs. The remainder of its liabilities comprise wholesale debt, which in the model are simply represented by bank bonds.

Chart 11 shows the resulting balance sheets in my simplified system. Netting out intra-financial system claims reduces the aggregate balance sheet from £42trn to £25trn. The 'HOT' sector has material holdings of UK bonds and equity and is in turn funded largely by issuing debt and equity. This will allow it to behave as a 'hot money' investor – driving the pricing of UK asset prices by constantly rebalancing its portfolio to take advantage of any deviation in yields across asset classes from its expectations of returns.

Chart 11: UK sector balance sheets

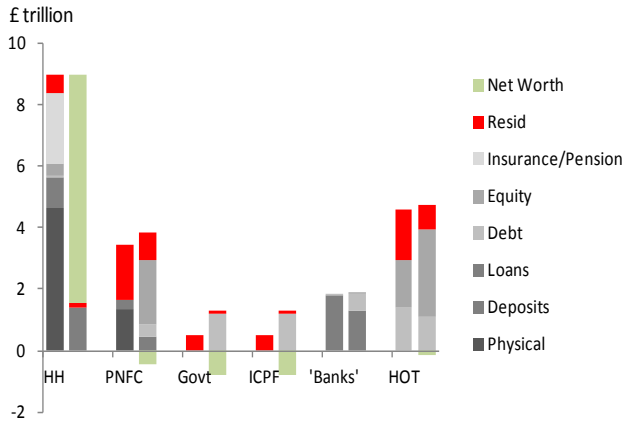

Source: ONS and Bank calculations

Aggregation of the output, expenditure and income data

In order to close the system, I need to reduce the granularity of the sectoral output, expenditure and income accounts. This is done in the top seven rows of Table B1 in Appendix B. Given that the focus of this model is to explore intra-sectoral financial flows, I have chosen a far greater simplification of the income flows than is generally used in macro models, as will become clear in the next section.

3.2.2 Behavioural relationships in the model

To transform the accounting framework described in Tables 1 and 2 into a model, one simply needs to add an equation to describe the behaviour of each of the variables in the system and equations for the price of each asset. In this section, I sketch out a baseline model which, in itself, does not allow me to explore financial fragility. Section 3.2.3 then discusses changes I have tried to make to the baseline model to capture some of the features of the crisis seen in the United Kingdom.

Income flows

In my various attempts at SFC modelling, I have been interested in finding an interesting role for financing flows and have thus kept the income, or 'real', side of the model extremely simple, in order to focus instead on intra-financial system flows. For example, I assume an economy that is entirely service based, which allows output to be produced on demand and allows me to abstract from accounting for inventories of goods, which adds a lot of complication to flow-of-funds models.[22] Households, firms, government and RoW all demand services, which are produced on demand by firms. Households consume as a function of their expected in-period income and their expected wealth. Expectations in the baseline model are simply set to their previous period value.

Government demand is also set exogenously. The government taxes the household sector a constant fraction of its wage income to pay for its expenditure, with debt issuance making up any shortfall. The rest of the income relationships are essentially assumed away, awaiting further development: firms' investment is set to equal their exogenous rate of depreciation, relegating capital to a trivial role in the model, with firms effectively producing solely out of labour; similarly, firms' imports are set to equal their exports, which in turn are set equal to the exports demanded exogenously by the RoW sector. This means the balance of payments for services is always in balance (and that exports and imports are again included to

[22] See, for example, Chapter Eight of Godley and Lavoie.

allow a richer treatment to be developed at a later stage). The insurance company and pension fund sector also takes a proportion of households' wages as a contribution to their long-term savings plans. For simplicity, the ICPF sector pays the same amount back to the household sector as in the form of pensions. The labour market also takes a very simple form, with households supplying labour perfectly elastically at a fixed wage and firms demanding sufficient labour to produce the amount of services demanded in that period.

The crude mechanics of the non-financial side of the model are very simple: aggregate output is pinned down by the exogenously determined government demand and a set of multipliers arising from the household's consumption function and the government's taxation of households. This leaves roles for financial flows to affect output through household's decisions to borrow to buy houses, through household wealth (which affects spending) and firms' financing structure (and thus their ability to produce output on demand).

The demand for and supply of assets

Table B lays out the assets and liabilities which need to be explained. This section goes through each sector's behavioural equations in turn.

Households: The only financial asset households hold directly in my model is deposits, with all their other asset holdings occurring through their passive subscription to ICPFs.

Non-financial firms: In my current version of the model, firms have a rather odd balance sheet. I fix their deposits and bonds exogenously and don't allow them to fluctuate. They issue or buy back equity to hit a leverage target, adjusting the supply of equity according partially each period to return to target. And they repay or increase bank loans each period to balance their budgets. This is artificial and unsatisfactory, but it leaves banks reliant on the availability of bank credit for production, as will be seen later.

Government: The government supplies sufficient debt each period to balance its budget constraint.

Banks: In the baseline model, banks passively accept deposits from and supply loans to households and corporates on demand. They purchase gilts to meet a liquid asset requirement and hold reserves to meet a capital requirement. Finally, they issue bonds to meet any gap between their assets and liabilities.

In an alternative closure, in which banks have insufficient capital to meet their capital requirement, they stop lending to firms and households until they re-build sufficient capital to lend again.

ICPFs: Insurance companies and pension funds manage assets on behalf of households. Their liability to households is not explicitly modelled as a stock: they levy a charge on households each period and make a pension payment to them each period. ICPFs manage a portfolio of assets. They rebalance this each period using a system of asset demand equations. Equation (1) shows Tobin's system of asset demand equations in matrix form. Multiplying through by ICPFs' expectation of their total portfolio size in period t gives their nominal demand for assets in that period, as shown in equation (2).

(1)

(2)

For ICPFs, the K assets in their portfolio are the bonds of banks, the government, NFCs and the hot money sector and the equity of banks, NFCs and the hot money sector. In the baseline model, ICPFs' expectations are assumed to adjust slowly from the return observed in the previous period to an exogenously determined long-run level, with the rate of adjustment governed by an adjustment parameter. This causes ICPFs to respond gradually to changes in the supply of assets or changing dividend yields and interest rates.

Hot money sector

The hot money sector faces a similar problem to ICPFs, in that it manages a portfolio of assets and does so using a system of asset demand equations, but because it incorporates the rest of the world, it also has liabilities to the UK, which in this model take the form of bonds and equity and are all held by the ICPF sector. Using the assumption that the UK is a small economy with a negligible impact on global asset prices, it is assumed that ICPFs can buy HOT bonds and equities with no impact on price.

The system of equations used by the HOT sector to determine asset demand looks the same as for the ICPF sector, except that it only includes five assets: the bonds of NFC, Govt and Bank and the equity of NFC and Bank. In the baseline model, expected returns are set equal to the yield on debt and the return on equity from a simple one-stage DDM for equity.[23]

The determination of asset prices

In contrast to the services and labour markets, prices clear asset markets. This gives supply and demand an explicit role in setting prices in the model. In each asset market, separate sectors' demands for the asset are first combined to calculate aggregate demand, using a simple set of accounting equations. Sectors' demands for each asset are determined as in equations (2) above. For domestic assets, prices are determined by dividing aggregate asset demand, which is a nominal variable, by aggregate asset supply, which is a real variable. For example, for NFC equity, the aggregate demand is the nominal value of NFC equity that ICPF and HOT want to hold, given their expectations of the return on equity; NFCs choose how

[23] i.e. return on equity = ———— , where divd = the dividend per share; P = the equity price; and g = expected growth (set exogenously here to 2.5%)

many shares to supply to the market; dividing the former by the latter gives a price per share. The prices of foreign assets are exogenous, given the assumption that the UK is negligibly small in global asset markets and thus a price taker.

Interest and dividend payments

The final set of equations required to close the system cover the mechanics of making interest and dividend payments between sectors. As explained above, bonds pay a coupon of 1 in each period. So the interest payment a bond-issuing sector must make is simply equal to the amount of bonds outstanding and the amount any particular sector should receive is simply equal to their holding. The interest paid on deposits and loans is the product of exogenously set interest rates and the amount of loans and deposits outstanding. In later versions of the model, bank interests could be determined endogenously, as in Godley and Lavoie Chapter 10. Finally, dividends paid by firms are calculated by applying a payout ratio to profits to determine the nominal value of dividends and then splitting it across shareholders according to their holdings.

3.2.3: Extensions to introduce financial crises

Adding a housing market

My first attempt at introducing financial fragility was to allow for a housing market. I assumed that the supply of housing was fixed (which is not that large an abstraction in the United Kingdom) and that old households sold a fixed fraction of housing to young households each period. The price young households were willing to pay was a function of their expectation of house prices, which in turn was an extrapolation of the previous period's house price gain. Banks would lend young households a multiple of their income to buy houses.

After an initial upwards shock to house prices, house prices would start to rise. Because I only had one aggregate household sector, the increased interest paid by households on their larger stock of debt was paid back to them in increased deposits. But their increased wealth did lead to an increase in spending and reduction in saving.

Multiple closures for banks and corporates

Without altering the model, house prices would continue to rise in this manner indefinitely. To get around this, I allowed the bank to switch between two types of behaviour. In normal times, it lent a higher multiple of income to households to buy houses and lent to corporate on demand. But if it had insufficient capital, it lent less. By lending less, young households had less leverage to buy houses, which depressed the price of houses (as a constant fraction was still sold and the price was set by dividing nominal demand by the real fraction of houses sold). This in turn triggered greater losses at banks, as write-offs on household lending were an asymmetric function of house-price changes. Furthermore, when banks

reduced lending to companies, this tightened their budget constraint, forcing them to produce less and pay down debt. Lower production meant lower employment and falling output. This cost further losses at banks, as household write-offs were a function of employment and corporate write-offs a function of output growth.

To date, I have not been able to get all features of the model working together. The reality imposed by using actual data as a starting point has been the toughest constraint.

4: How all this relates to the problems with finance in macroeconomics

I think the work I have sketched out here gives me some hope in dealing with the failings of mainstream macro models to help me in my work.

Not including credit for asset purchases and ignoring banks' role in credit creation: I think systematic analysis of the FoF data can go a long way to getting around this. Making sure that the stories we tell to understand economic develops cover balance sheets and asset prices and do not simply extract from them is a good start. Going beyond that and explaining it with models will be harder, but I think SFC models that incorporate sectoral balance sheets offer a lot of hope.

Equilibrium: This is easily dealt with in the data work by constantly remaining open to the idea that the economy is evolving, rather than being shocked away from a state of equilibrium. I think it can be dealt with explicitly in the SFC modelling approach by allowing multiple closures of the model – one for a buoyant world of rising asset price expectations and easy credit and one for a world of falling asset price expectations and contracting credit. The expansionary world will not have stable equilibria (in the sense that debt/GDP is constant). Although such models can say little about what tips systems from expansion to contraction, they can probably be useful for pointing out that certain types of behaviour are unsustainable.

Representative agents: This is not satisfactorily dealt with in either the FoF data approach or the subsequent SFC modelling. The data approach needs to be augmented with micro data on the underlying agents that comprise each sector. We are attempting to do this at the Bank by building a database that allows us to build up to sectoral balance sheets bottom-up from firm level balance sheets. This is easy for banks, for which we have regulatory data on the whole sector, but harder for firms, where we have to rely on much poorer data and grapple with issues of residency versus ownership when comparing published accounts to the national accounts.

In the SFC modelling, I have tried to get around this by splitting the household sector into those with different types of balance sheet – rich, old households that own most of the economy's assets; middle-aged, middle income households that have mortgages and only hold financial savings through pensions;

and younger and poorer households that just have very modest deposit savings. So far, the detail has become too complicated to get the housing market to work.

The wrong sort of micro-foundations: The SFC modelling approach I have taken is explicitly sector-based rather than agent based and so completely fails this criticism at this point. Along with a colleague at the Bank, I have begun talking to some external researchers about implementing the model in an agent-based framework, in the hope that we can first nest a simplified balance sheet model, then relax the assumptions of homogenous agents in each sector and gradually introduce more interesting behaviour, while maintaining the rigour of an accounting framework.

Bibliography

Barwell, Richard and Burrows, Oliver, Growing Fragilities? Balance Sheets in the Great Moderation (April 12, 2011). Bank of England Financial Stability Paper No. 10. Available at SSRN: http://ssrn.com/abstract=1815000

Gai, Prasanna and Kapadia, Sujit, Contagion in Financial Networks (March 23, 2010). Bank of England Working Paper No. 383. Available at SSRN: http://ssrn.com/abstract=1577043

Galati, Gabriele and Moessner, Richhild, Macroprudential Policy - A Literature Review (February 1, 2011). BIS Working Paper No. 337. Available at SSRN: http://ssrn.com/abstract=1858689

Gertler, Mark & Karadi, Peter, 2011. "A model of unconventional monetary policy," Journal of Monetary Economics, Elsevier, vol. 58(1), pages 17-34, January.

Godley, W and Lavoie, M (2007), Monetary Economics: An integrated approach to credit, money, income, production and wealth, Palgrave/Macmillan.

Jacome and Nier, IMF, "Macroprudential Policy: Protecting the Whole"
http://www.imf.org/external/pubs/ft/fandd/basics/macropru.htm

Haldane, Andrew G, April 2009, "Rethinking the financial network"
http://www.bankofengland.co.uk/publications/Documents/speeches/2009/speech386.pdf

Starr-McCluer, M (1998), 'Stock market wealth and consumer spending', Federal Reserve Board of Governors Finance and Economics Discussion Series, 1998-20.

Appendix 1: An accounting framework for stock-flow consistent modelling

Godley and Lavoie (2007) build a series of closed accounting frameworks, which encompass the standard national income flows, such as wages and consumption, the counterpart financing flows, such as bank loans and deposits and stocks of physical and financial assets and liabilities. This framework lends itself to representation in a set of matrices. The first matrix captures flow variables (Table 1a). The columns represent the sectors of the economy and the rows represent the markets in which they interact. The top half of the table covers the bread and butter of mainstream macroeconomic models – the standard income, expenditure and production flows associated with the national accounts – while the bottom half covers sector's financing flows. The matrix has two important properties. Each sector's resources and uses columns provide their budget constraint – they sums must equal to ensure that all funds they receive are accounted for. And each row must also sum to zero, to ensure that each market clears – that is, the supply of a particular asset must be matched by purchases of that asset, to ensure that no funds go astray. Finally, note that the table can usefully be split in two, with the top half covering the standard income and expenditure flows and the bottom half covering financing flows. The two halves of the table are linked together by each sector's 'net lending balance', or 'financial surplus'. The net lending balance can be used to summarise each sector's income and expenditure flows as the difference between the amount the sector spends on consumption and physical investment and the amount that it receives in income. This difference must be met by financing flows – either borrowing or the sale of financial assets. In national accounts terminology, a sector's net lending balance (NL) must equal its net acquisition of financial assets (NAFA) less its net acquisition of liabilities (NAFL). Across sectors, the net lending balances have to sum to zero, as all funds borrowed by one sector must ultimately come from another.

While it is useful to split the table for accounting purposes into income and expenditure flows and financing flows, it is important to note that the acquisition of financial assets and liabilities is not necessarily determined purely by imbalances between income and desired expenditure. Sectoral balance sheets can adjust for other reasons. Agents may want to borrow money to purchase assets, simultaneously acquiring financial assets and liabilities. And on occasion agents may want to shrink the size of their balance sheets, selling off financial assets to pay off financial liabilities. Finally, some agents may default on their debt obligations, which will involve a revision in the financial assets and liabilities of both debtor and creditor. At an aggregate level, simultaneous expansion of a sector's assets and liabilities invariably represents one set of underlying agents taking on assets whilst the other takes on liabilities. The household sector provides an important example. If a young household takes a mortgage to buy a house from an old household, the sector in aggregate simultaneously acquires a liability (the young household's mortgage) and an asset (the deposit created for the young household to pay to the old household).

All of these activities – leveraging up, deleveraging and default – involve NAFA and NAFL moving in lockstep. The net lending identity still holds: the gap between income and expenditure determines the

difference between NAFA and NAFL. But the absolute size of the NAFA and NAFL flows is determined by agents' actions in financial markets.

Table 1.a: stylised flow matrix

		Households		Firms		Banks		...	RoW		∑
		resources	*uses*	*resources*	*uses*	*resources*	*uses*		*resources*	*uses*	
income flows	Wages	+ wages			- wages						0
	Consumption		- goods	+ goods							0
	Interest payments/receipts	+ interest			- interest	- interest	+ interest				0
	...										0
	Net lending	$resources_{hh}$ + $uses_{hh}$ = NL_{hh}		$resources_{firm}$ + $uses_{firm}$ = NL_{firm}		$resources_{bank}$ + $uses_{bank}$ = NL_{bank}					0
		resources = NAFL	*uses = NAFA*	*resources = NAFL*	*uses = NAFA*	*resources = NAFL*	*uses = NAFA*				
financing flows	Δ Loans			+ loans			- loans				0
	Δ Deposits		- deposits			+ deposits					0
	...										0
	∑	$NAFA_{hh}$ = NL_{hh} + $NAFL_{hh}$		$NAFA_{firm}$ = NL_{firm} + $NAFL_{firm}$		$NAFA_{bank}$ = NL_{bank} + $NAFL_{bank}$					

1.b: stylised balance sheet matrix

	Households		Firms		Banks		...	RoW		∑
	assets	*liabilities*	*assets*	*liabilities*	*assets*	*liabilities*		*assets*	*liabilities*	
loans				- loans	+ loans					0
deposits	+ deposits					- deposits				0
domestic equity	+ equity			- equity				+ equity		0
physical capital			+ capital							capital
...										
Net worth		(= L_{hh} - A_{hh})		(= L_{firm} - A_{firm})		(= L_{bank} - A_{bank})			(= L_{RoW} - A_{RoW})	- capital
∑		0		0		0			0	0

The second table (Table 1b) captures the balance sheet positions of each sector. The balance sheet matrix is updated over time using data on the acquisition of assets and liabilities from the transaction flows matrix, and revaluation effects to asset positions. Proceeding in this manner, balance sheets always balance across sectors, flows of funds are always accounted for over time and the impact of flows of funds on balance sheets is always recorded.

The design of such a framework entails interesting questions about the degree of granularity to introduce. Financial fragility tends to lurk in the tails of weak institutions within any particular sector, rather than being uniformly distributed, which argues in favour of more granularity. Similarly, the degree of disaggregation of instruments is of interest: should we aggregate up all fixed income instruments and call them debt, or do we want to break them down along maturity lines to try to identify maturity mismatch? In practise, the answers to these questions are pragmatic, determined by the availability of data.

Appendix B

Table B1: Balance sheet matrix

	HH	NFC	Govt	Domestic Banks	ICPF	Hot Money	∑
Deposits							
HH	+ dep_hh_d			- dep_hh_s			0
NFC		+ dep_nfc_d		- dep_nfc_s			0
Loans							
HH	- loan_hh_d			+ loan_hh_s			0
NFC		- loan_nfc_d		+ loan_nfc_s			0
Bonds							
NFC		- debt_nfc_s			+ debt_nfc_h__ICPF	debt_nfc_h__HOT(i),	0
Govt			- debt_govt_s	+ debt_govt_h__bank	+ debt_govt_h__ICPF	+ debt_govt_h__HOT	0
Bank				- debt_bank_s	+ debt_bank_h__ICPF	+ debt_bank_h__HOT	0
HOT					+ debt_hot_h__ICPF	- debt_hot_s	0
Equity							
NFC		- equity_nfc_s			+ equity_nfc_h__ICPF	+ equity_nfc_h__HOT	0
Bank				- equity_bank_s	+ equity_bank_h__ICPF	+ equity_bank_h__HOT	0
HOT					+ equity_hot_h__ICPF	- equity_hot_s	0
Insurance/pension claims	+ pension_d				- pension_s		
Bank reserves				-reserves			
∑	net wealth	net wealth	net wealth	net wealth	net wealth	net wealth	0

Table B2: transaction flow matrix

		HH	NFC		Govt	Domestic banks		ICPF	Hot money	
		current	current	capital		current	capital			∑
Income flows	G		+ G_s		- G_d					0
	C	- C_d	+ C_s							0
	I		+ I_s	- I_d						0
	M		- M_d						+ M_s	0
	X		+ X_s						- X_d	0
	Lab	+ W * L_s	- W * L_d							0
	T	- T_s			+ T_d					0
Interest and dividend payments	Pension contributi	- pencontr_s						+ pencontr_d		0
	Interest on:									0
	deposits					*interest * debt* [1]				0
	loans									0
	bonds									0
	Dividends					*dividend * equity* [2]				0
Δ assets	Δ assets									
	deposits					*- Δ assets* [3]				0
	loans									0
	bonds									0
	equity									0
	pension claims									0
	bank reserves									0
	Pension payout	+ penpayout_d						- penpayout_s		0
	Write-offs	+ wor_hh * loan_hh_s(-1)				- wor_hh * loan_hh_s(-1)				0
			+ wor_nfc * loan_nfc_s(-1)			- wor_nfc * loan_nfc_s(-1)				0
	∑	Δ *Net wealth*	Δ *Net wealth*	Δ *Net wealth*	Δ *Net wealth*	Δ *Net wealth*	Δ *Net wealth*	Δ *Net wealth*	Δ *Net wealth*	0

For the sake of succinctness, some types of relationship have been described as a class in this table, rather than writing out each relationship in full:

(1) all holdings of debt, as laid out in Table B, entail an interest payment from the liability issuer to the asset holders

(2) all holdings of equity, as laid out in Table B, entail a dividend payment from the liability issuer to the asset holders

(3) all asset holdings, as laid out in Table B, can change in each period