Speech

Household behaviour and policy analysis

Speech given by

Martin Weale, External Member of the Monetary Policy Committee

New Zealand Government Economists' Network 2nd Annual Conference, Wellington

14 December 2012

Introduction

Thank you very much for inviting me here to talk to this conference. It is, of course, a great honour and I should also say how much I am enjoying my first visit to New Zealand. At the same time preparing this talk has been something of a challenge. I am very aware of that there is a wide range of interests represented here in this room and I can address only a few of these.

I would like to talk to you about economic modelling and policy analysis with a particular reference to limitations of analysis based round the concept of a representative agent. I will do this by focusing on models which represent the economy as a collection of individual households in different circumstances and their use to address important policy questions. Even if they do not provide all the answers to the questions we might wish to ask, they have the prospect of being able to help policy-makers come to informed decisions and understand the pressures that they face. I would like to describe the way in which such models can be used to address a heterogeneous range of topics: the effects of credit constraints and fears about credit availability, tax structure, social security and pension arrangements, and influences on the take-up of education by mature students. I will also say something about one aspect of behavioural economics, myopic behaviour. We are often told that a short-coming of "traditional" economics is that it fails to take account of this last issue. Does it really matter? But models of household behaviour can never be models of the whole economy. I would like to end my talk by making a brief reference to a study which has put the structures I describe into a general equilibrium framework. This makes it possible to look at the impact of monetary policy on people in different circumstances, a topic of considerable importance in the United Kingdom today.

What is the difference between the approach I am describing and the models generally used by macro-economists? The latter generally rely, to a greater or lesser extent, on the assumption that the economy can be modelled with reference to the behaviour of representative households and representative firms. If households can be assumed to have infinite horizons and do not face wealth constraints and individual risks can be fully insured, then it is not necessary to distinguish young people from old people or to worry about the consequences of a situation where some people have debts and others have substantial holdings of assets. Similarly, if there is effectively a representative firm, it is not necessary to worry about whether the behaviour of businesses might change if, perhaps as a result of particularly adverse demand shocks, they have become more indebted than the average.

The key importance in moving beyond representative firms and households is that it makes it possible to explore the ways in which households' and firms' behaviour may be affected by their circumstances, even if everyone has the same preferences and technology. That may be of interest in its own right. But if such effects are asymmetric, then they will also have implications for the evolution of the whole economy. Whether and when such effects are likely to be important can be known only by examining them.

Constrained Households

Let me give a simple example of the way in which models of this type can differ in important respects from traditional models based on a representative agent; this example follows Deaton (1991). I consider a situation in which there is no economic growth and where the rate of interest is equal to the rate of discount. In such a world, a representative consumer will smooth out the effects of fluctuations in income; consumption will be set at a level which it is expected can be maintained indefinitely. In such a situation most analysis of the effect of risk suggests that it has relatively little impact on this conclusion.

But now suppose that people face borrowing constraints. If their incomes are expected to rise over the life cycle, then they will naturally want to borrow when young. You might expect that, in such circumstances, young people will consume the whole of their incomes and start saving only after income has risen somewhat. Indeed the traditional model is often extended by making the assumption that some proportion of income goes to people who spend all of it. This is justified on the assumption that they would like to be indebted, or more indebted than they are, but that their access to credit is constrained.

If, however, I simulate a panel of people whose incomes are volatile, then there are always some who have had a run of bad luck and find themselves firmly against a borrowing constraint. But the existence of the borrowing constraint much increases the importance of precautionary saving. The average young person, at least in the absence of social security benefits, will typically spend less than their income when young. Consumption will rise steadily and then level off when people typically feel they have built up enough balances to protect themselves adequately from the effects of a run of bad luck on their spending power. If they have to make their own provision for retirement, they will also save up to finance consumption in old age, after taking account of any state pension. In the example shown below I have assumed that people work for forty years, and live for another twenty, during which they receive a modest state pension. But borrowing is not possible. I show mean consumption as a function of age; you can see that consumption rises and then levels off as expected.

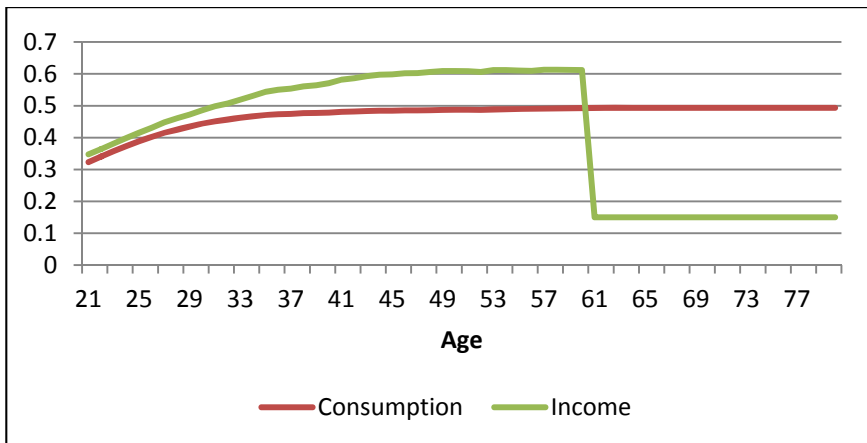**Chart 1: Mean Consumption and Income by Age**

Figure 1: Consumption and income by age when Borrowing is not Permitted

Now suppose that people can borrow an amount up to 0.3 units but that debts have to be repaid before they reach retirement. Facing less of a constraint, young people now go substantially into debt, expecting to repay their debts out of rising incomes. Chart 2 shows the impact on mean consumption measured relative to the path of

Chart 1.

Chart 2: The Impact on Consumption of allowing Borrowing
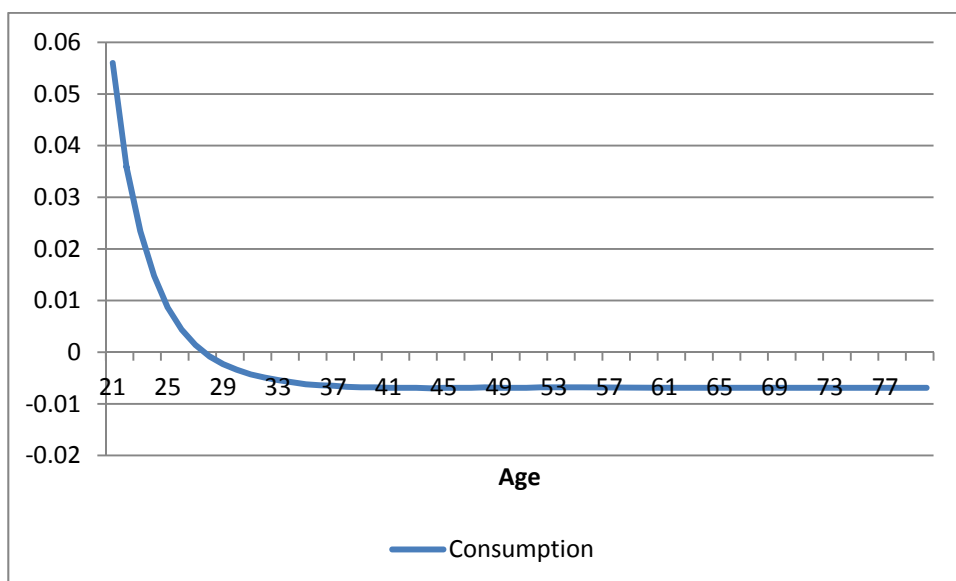
As you can see people take advantage of the extra capacity for borrowing. Consumption by young people is appreciably higher, but at the cost being lower later in life. Access to credit makes it easier for people to smooth their consumption and should be expected lead to a higher level of life-time welfare. Indeed, the calculations underlying this simple and purely illustrative framework suggest that a lump-sum payment of around 10% of initial income would be needed to compensate a twenty-year old on average income at that age for the life-time absence of access to credit. Whether this sort of sum should be regarded as small or large depends on the eye of the beholder. And the estimated cost is in any case based on a highly stylised model.

Such calculations are enough to remind us that we should not lose sight of the costs associated with restrictions of this type. It may be a good idea to limit access to credit but one should certainly not assume it is costless. While I am not involved in the macro-prudential side of the Bank of England's work, I can understand why the Financial Policy Committee has not, as yet, requested the power to set credit limits as part of their macro-prudential tool-kit.
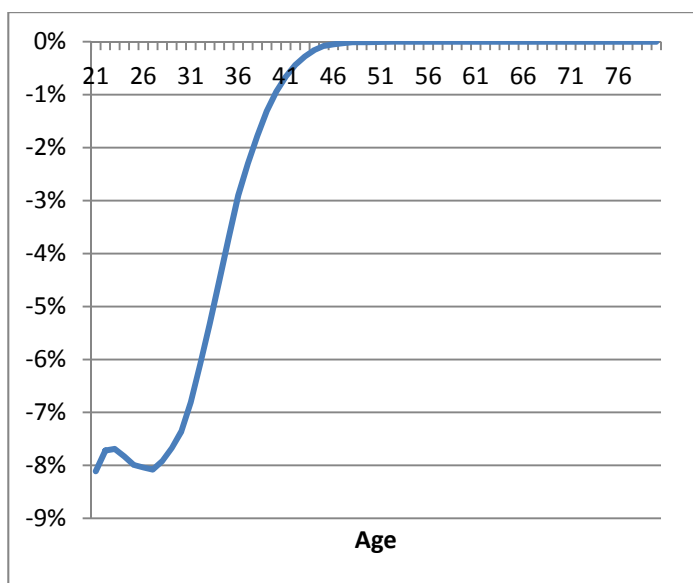
Fear of a Credit Crunch

If people expect to be able to borrow some will do so, and others will accumulate less wealth than they would in the absence of the ability to borrow. But now suppose that, having borrowed, people fear they will be asked to repay their debts, and that they will, in future not be able to borrow. This fear will itself trigger a reduction in consumption.

I can illustrate this by making the assumption that the chance people put on this prospect of a credit crunch increases from 10% to 80% for some exogenous reason. The impact of this on consumption cannot be derived from optimisation alone. It also depends how rapidly those who are at risk of becoming excess borrowers think they will have to pay down their debts. I have assumed that a credit crunch means i) debt limits are reduced to zero and ii) people have to devote up to half of their incomes to clearing debts, subject to a minimum consumption allowance set at 1/3 mean income. With these assumptions an increase in the risk of a credit crunch from 10% to 80% triggers a reduction in aggregate consumption of 1.7%.

Chart 3 shows the immediate effect on the level of aggregate consumption measured relative to that observed with only a 10% risk of a credit crunch. This wears off gradually. But I am not sure that the long-run impact is of great interest since it is unlikely that consumers would go on giving a high probability to a credit crunch if that did not eventually happen.

Chart 3: The Impact on Consumption of Fear of a Credit Crunch

Source: Bank simulations

Chart 4: UK Experience since the Crisis
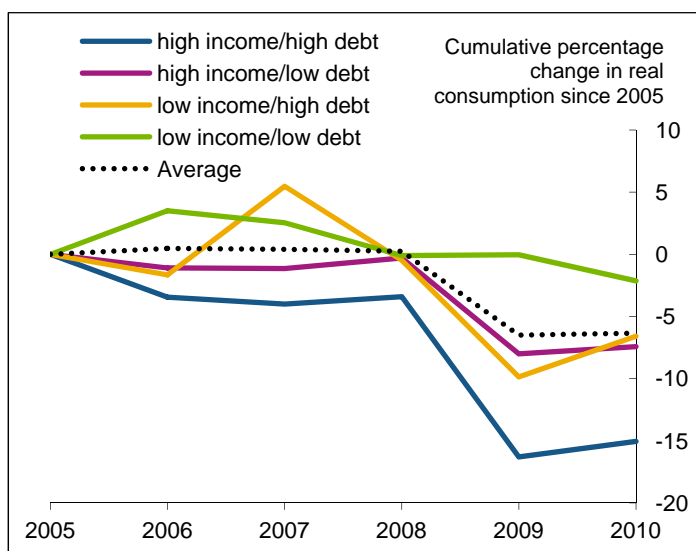
I am not sure whether Chart 3 tells us that fear of a credit crunch can have a significant impact on consumption or whether we should conclude that it will do so only if people fear draconian terms for debt repayment[1]. In the United Kingdom, data from the Living Costs and Food Survey (see Chart 4) certainly suggest that the consumption of high income high debt people has fallen markedly since the crisis. But

---

[1] Alan, Crossley and Low (2012) identify weaker effects, because they look at the impact of a restriction on new credit rather than fear of rapid repayment of any existing debt.

plainly rather more work is needed to see whether or how far it can be explained by the phenomenon that I have described.

One more general point follows from this, however. We often hear that households are deleveraging. This calculation suggests that people are more likely to be tolerant of any given level of debt it they believe that they can obtain further credit, should the need arise, than if they think they are up against a credit limit. So a fear that credit is tight itself encourages deleveraging.

While I will return to the macro-economy at the end of this talk, I would now like to provide some other illustrations of the way in which modelling of this type can be used to address economic policy issues.

Capital Income Taxation

Perhaps the most obvious use of disaggregate modelling of households is in the study of tax and benefit policy. Indeed this presents a very good example of a situation where a macro-economic model provides a stark conclusion which is, perhaps, not very robust. Should income from capital be taxed? Macro-economists have argued that it should not be. Thus Lucas (1990), following Chamley (1986), wrote:

"When I left graduate school, in 1963, I believed that the single most desirable change in the U.S. tax structure would be the taxation of capital gains as ordinary income. I now believe that neither capital gains nor any of the income from capital should be taxed at all."

So why did Lucas believe this in 1990? In the traditional framework people accumulate capital to the point where the marginal product of capital, net of depreciation, balances the rate of time preference. A tax on income from capital reduces the return on capital and results in the economy becoming less capital intensive, so as to restore the previous relationship between the discount rate and the post-tax rate of return. In effect the return on capital is left unchanged, but labour incomes are reduced because labour has less capital to work with.

Conesa, Kitao and Krueger (2009) explore this issue in a simulation model designed to represent, in broad terms, the circumstances of American households. They assume the policy aim is to maximise the average expected life-time welfare of young people, using the approach first developed by Mirrlees (1971). They find, not surprisingly, that the optimal outcome depends on what they take as a given and, in particular on the way in which labour supply responds to changes in the real wage. Borrowing constraints are, despite my earlier analysis, not particularly important and nor is the fact that people are exposed to uncertainty about their earnings against which they cannot insure. Their results, with flexible labour supply, are summarised in Table 1.

Seen in purely aggregate terms the results suggest that their optimal tax structure is bad for "the economy". Output, labour supply, the capital stock and consumption are all reduced. The results that they show are driven by the fact that their policy affects the distribution of life-time resources. In turn this demonstrates the general point, that desirable outcomes cannot be identified independently of society's preferences. Lucas' observations were conditional on the social preferences embodied in his use of a representative consumer. Models, whether aggregate or disaggregate do not tell us what social preferences should be and I am certainly not offering any view on that, any more than I am offering a view on desirable tax structures. Nor do they tell us about the international issues associated particularly with corporate taxation. But the point I want to stress is that a simple macro-economic model based on a representative consumer may not deliver robust conclusions.

Table 1: Optimal Income Tax Structure for the United States suggested by Conesa *et al.*

Capital Tax Rate	36%
Labour Tax Rate	23%
Impact relative to benchmark economy	
Average Hours	-0.56%
Total Labour Supply	-0.11%
Capital Stock	-6.64%
Output	-2.51%
Aggregate Cons.	-1.63%
Cons Equiv. Var	1.33%
Of Which	
Cons. Total	1.29%
Level	-1.63%
Distribution	2.97%
Leisure. Total	0.04
Level	0.41
Distribution	-0.37

Source: Conesa *et al.* (2009)

Application to Retirement Benefits

Disaggregate models can also be used to explore other aspects of taxes and transfers. For example Sefton and Van de Ven (2009) use a similar approach to explore the design of retirement (or strictly old-age) benefits in the United Kingdom. They find that, with their particular model structure, the solution is not clearly defined, so that changes in the taper rate applied to benefits have, over a fairly large range, relatively little impact on the expected life-time welfare of the median household at the start of its working life. This suggests, in turn that the results from exercises of this type may be sensitive to the specification of the

models used. That is hardly the fault of the models or of the technique but it is something that users of such results need to be aware of.

Other examples where modelling the effects of the sort of random events that can affect people over their life-course are easy to think of. But extensions require the introduction of extra variables and one lesson from working with such models is that not many straws are needed to break the camel's back. Nevertheless, in principle they could be used to explore issues such as the financing of long-term care for people who need it.

Investment in Human Capital

I would now like to turn to a rather different topic, that of education. Many countries have tried to increase the proportion of school-leavers going to university or to other forms of education and training. In Britain participation in tertiary education, measured here as the proportion of nineteen year-olds receiving it, increased from 30% in 1985 to 36% in 2010[2]. Efforts have also been made to increase greatly the opportunities to gain qualifications for those who did not go on to tertiary education immediately after leaving school. Thus in the United Kingdom the number of mature university students has increased sharply between 1990 and 2010. But, underlying this, there is something of a puzzle.

Mcintosh (2006), working with data from the Labour Force Survey, suggests that compared to someone who leaves school at 16 with five good subject passes at GCSE, someone who stays at school and then completes a university degree can, on average, expect to earn about 40% more. Moreover there is good evidence that participation in higher education is sensitive to factors such as social class and parental educational attainment. It is rather unlikely that these factors are closely correlated with ability to benefit, and it is therefore more probable that, despite the sharp increase in participation which has taken place, there remains a large pool of people who would benefit financially from higher education were they to embark on it. But they have not done so.

The opportunity to go on to university immediately after school, if missed, does not come back. But those who missed out can certainly participate in higher education[3] later in working life. Despite the earlier observations on numbers of mature students, data are scanty about the overall pattern of educational upgrading to gain higher education qualifications. The British Household Panel Survey suggests that about 5% of men educated to a level slightly higher than five good GCSE passes (Level 2) do so after the age of twenty-five, as do just under 20% of men with at least two A-levels or equivalent (Level 3).

Once again, it appears that there are substantial benefits to upgrading, if smaller than those to a degree taken after a university course following on immediately from school. Van de Ven and Weale (2012)

---

[2] OECD data
[3] It should be noted that higher education includes study for sub-degree qualifications such as Higher National Certificate and Higher National Diploma, as well as degree courses.

suggested that a man who is qualified to a level slightly higher than GCSEs gains an average increment to their earnings of around 20%, while someone with better initial qualifications gains an increment of, on average 9%. But there are several sources of uncertainty. First of all, like any parameters, the estimates of the average benefit are not precisely determined. Indeed the figure of 9% is significant at a 10% level but not at a 5% level. But secondly, and much more importantly, these are population averages; analysis of the data contained in the British Household Panel Survey puts the standard deviation of the individual effects at over 15%. So, as with much of life, there is no cast iron guarantee.

In addition, of course, embarking on a course does not mean you will complete it successfully. I have not been able to find any statistics on failure rates, and it is clear why household surveys ask people about exams they have passed, and not also about exams they have failed. But universities report drop-out rates of over 30% among mature students; there must be a risk that anyone embarking on a course will find that he or she cannot complete it for one reason or another.

Of course, risk is not the only factor offsetting the benefits expected from study. There are direct costs involved, both in terms of fees and in terms of the time taken up in study. In the United Kingdom university fees, having been abolished in the 1970s, were re-introduced at £1000 p.a in 1998. After an earlier increase to £3000, universities have most recently been allowed to raise fees to up to £9000. At the same time it has to be remembered that many sub-degree higher education courses have much lower fees. And they take less than the traditional three years of full-time study. In the simulations we have described here, it has been assumed that mature students can either study full-time and gain a qualification in two years, or study part-time and gain it in three years. In the study we also assumed that it is not possible both to study full time and to work full time. But full-time students can work part-time should they wish.

The simulation model offers an ideal means of exploring how far the direct costs of higher education and the uncertainty about its outcomes can account for the prevalence of upgrading summarised above. Expected welfare can be calculated both with and without a higher education qualification and, at each point, individuals can be assumed to choose whichever path delivers the higher level of expected welfare.

Preliminary results presented in Tables 2 and 3 show the life-time prevalence of upgrading generated for different fee rates and different subjective risks of failure, and I think the conclusions from these are quite stark. They suggest that, even if one allows for the observed uncertainty surrounding the benefits of adult education, it is hard to explain, in terms of plausible observed costs[4] why the take-up rate is so low. I am not saying that subjective risk of failure is necessarily the explanation of this. But the figures do show that some other factor is needed to explain the prevalence. Whether this is fear of failure, ignorance of the benefits or inertia, some high subjective cost of making a decision, is something which cannot be answered by a study

---

[4] With part-time study assumed to cost half of full-time study.

of this type; it requires some sort of detailed study looking into the reasons why people are unwilling to embark on what seems to be a sound investment.

Table 2: The Proportion of Men upgrading from Level 2 to Level 4 (per cent)

| | Risk of Failure | | | | |
Cost of full-time course (£ p.a.)	0%	20%	40%	60%	80%
0	81%	75%	64%	46%	11%
500	80%	73%	61%	43%	8%
1000	79%	71%	60%	39%	5%
3000	73%	64%	51%	28%	0%

Observed prevalence 5%

Source: Van de Ven and Weale (2012)

Table 3: The Proportion of Men upgrading from Level 3 to Level 4 (per cent)

| | Risk of Failure | | | | |
Cost of full-time course (£ p.a.)	0%	20%	40%	60%	80%
0	71%	62%	52%	37%	7%
500	69%	60%	50%	34%	5%
1000	67%	59%	48%	31%	3%
3000	62%	54%	42%	23%	0%

Observed prevalence 18%

Source: Van de Ven and Weale (2012)

Two important points follow from these simulations. First of all, uncertainty about the returns to education is a factor influencing take-up. Without it, the participation rates would be close to 100%. But given this, fees are not on their own enough to account for observed participation. However, a subjective risk of failure could be a substantial reason for not wanting to attempt to gain a higher education qualification; since, if this risk were to rise to 100%, there would be no participation, it is plain that it could account for the observed participation rates. Nevertheless, as the table above shows, this perceived risk has to be high before participation rates are brought down to those observed, even if full-time fees are in the rate of £1000-£3000 p.a. as they probably were in the period to which the data relate.

These results certainly suggest that participation in higher education by mature students is unlikely to be very sensitive to the fees charged, at least over the range considered. But, once again, before drawing any firm conclusions, it would be necessary to examine the robustness of the findings above to a range of behavioural assumptions.

Myopic Behaviour

Let me now move on to a rather different issue. It is sometimes argued that one of the shortcomings of traditional economic analysis is that it does not reflect the lessons of behavioural economics. There are a number of ways in which behavioural economic individuals are thought to differ from rational economic individuals; I would like to focus on the idea that people worry rather little about the near future relative to the present. They would not pay very much for £1 in a year's time relative to £1 today, but would then pay only a little less if forced to wait a further year for their £1. It is commonly thought that this might disturb conventional assumptions about people's choices, as indeed it might. This behaviour is often described as myopic.

Such behaviour is quite different from simply assuming that people have high discount rates. An individual with a high discount rate will spend their money early rather than late in life but will not regret their choice. An individual who discounts the near future excessively will regret their decisions later on, and, if they understand their behaviour, may take steps to prevent themselves from making choices they will subsequently regret (see Diamond and Köszegi, 2003) – rather like asking someone else to keep the key to the drinks cupboard because otherwise it is difficult to avoid the temptation to over-indulge.

For people who see themselves as prone to over-spending, one obvious means of locking up one's savings is to join a pension scheme. It has its down-side, of course, because money invested in a pension scheme is not available to meet temporary needs for cash before retirement. So we can see that it is by no means obvious that a myopic individual will put less than a rational individual into a pension scheme. Nevertheless, by examining participation in pension schemes in a fully articulated model of the type used to generate the earlier results, Van de Ven (2010) was able to estimate the degree of excess discounting needed to account for pension participation in the United Kingdom. He found a discount factor of 0.85 as compared to the discount factor for subsequent years of 0.98. This degree of excess discounting is, nevertheless much less marked than the value of 0.68 found by Laibson *et al.* (2007) or 0.30 reported by Fang and Silverman (2009) on data for the United States.

These differences may arise because the latter studies considered only very specific subgroups of the population, or perhaps because they were carried out in a framework in which labour supply was assumed fixed, and so it was much harder to "make up" the effects of early over-spending. It should be noted, however, that while myopia improved the capacity of the model to explain pension participation, it worsened its ability to explain consumption. Or possibly, as Diamond (2008) suggests, that value of the excess discounting parameter may be context-specific. In any case, it would be quite wrong to assume that, if evidence for myopia is found, that will necessarily make it easier to understand people's spending decisions. Table 4 shows simulated average pension wealth held by households with heads aged 35-49 as a multiple of their median annual income. The table shows how this depends on the return on pension saving and on the excess discount factor. It shows one cannot conclude that greater myopia always leads to lower saving.

Table 4: The Effects of Excess Discounting (Myopia) on Pension Saving

Pension Return	Excess Discount Factor		
	0.7	0.85	1
2.5% p.a.	0.86	1.01	0.85
4% p.a.	1.51	1.72	1.74

The table shows pension wealth held by households aged 35-49 as a multiple of median income

Source: Van de Ven 2012

I remain to be convinced that myopia of this form is an important influence on people's behaviour. I would now like to return to the macro-economy.

Towards a General Equilibrium

The analyses I have presented have all been partial. They have not addressed the behaviour of firms and they have taken prices as fixed. For this sort of approach to be used more widely by macroeconomists it needs to describe in full the workings of the economy. Gornemann, Kuester and Nakajima (2012) present what I believe to be the first attempt to merge models of heterogeneous agents with a New Keynesian production and pricing structure so as to examine the distributional impacts of monetary policy[5].

Such an issue is of great interest to me. In the United Kingdom the policy of low Bank Rate followed by asset purchases which have reduced yields on government stock has led to complaints from those people, typically retired, who supplement their pensions from interest income. People close to retirement, who are thinking of buying annuities (which is more or less compulsory in the United Kingdom) are also unhappy and pension funds have noticed deficits opening up as the yields they use to discount their liabilities have fallen. As Bean (2012), Miles (2012) and Bank of England (2012) make clear, the policies pursued by the Bank have also had the effect of raising the prices of both gilts and equities; the implications of this also need to be factored in to the calculations. But any observer of the debate would note that the use of monetary policy is not as politically simple as Woodford (2011) suggests:

"to the extent that the problem [insufficient nominal expenditure] can be solved using monetary policy, it is costless to do so, since monetary policy has no other aims to fulfil".

Anyway Gorneman *et al.* find that, in an environment where there is only one financial asset, tight money favours wealthy (and typically old) people while easy money presumably has the opposite effect. Table 5 summarises the effects of a positive productivity shock and also unexpected monetary tightening. Of course, with only one financial asset this model does not show the capital gains and losses which may result from

---

[5] Oh and Reis(2011) look at the effects of transfer payments and fiscal policy.

monetary policy shocks; these can have important distributional consequences, but they are also less apparent, if not always less important, than the income effects of monetary policy. Table 5 certainly suggests that tight money is popular with savers with implication that, as we have found in the United Kingdom, low interest rates are not. The point is not that Table 5 tells us something we did not know beforehand but rather that it illustrates an effect which is left out of traditional analysis.

Table 5: The Distributional Effects of a Productivity Shock and a Monetary Policy Shock

% change in life-time consumption equivalent	Productivity Shock (+1 % point for 1 year)	Monetary Policy Shock (+ 1% point for 1 year)
Representative Agent	0.109	-0.014
Average of Households	0.151	-0.05
By Wealth		
Top 5 per cent	0.13	0.022
40-60 per cent	0.122	-0.034
Bottom 5 per cent	0.168	-0.104

The table shows the effect of temporary shocks valued in terms of their equivalent in life-time consumption.
Source: Gornemann *et al.* (2012).

As I noted above, the picture offered here is only a part of the story because the effects of capital gains or losses also need to be taken into account; these have been an important feature of the bank's asset purchases and, as Bank of England (2012) explains, for many savers they may have more than offset the flow effects described above. But the more general point to consider is whether an analysis of the type set out above would, if extended to take account of these gains and losses[6], help in our understanding of the effects of monetary policy. There are two reasons why it should. First of all, it is natural to think that the direct impact of interest rate changes depends in part on the balance between savers and borrowers in the economy; this is likely to change over time. Secondly, a framework which takes account of people's individual circumstances is likely to be a better basis for showing the effects of capital gains and losses than is one which relies only on a representative agent.

At the same time we need to remember that such an approach would be subject to the other standard criticisms of the New Keynesian framework in its current incarnation. Goodhart (2007) and others have pointed to the lack of any financial structure or role for money. Plosser (2012) has summarised concerns about the assumptions on price adjustment. And my guess is that the model of firm behaviour described by Bloom (2009), with its sensitivity to uncertainty, its analysis of the option to invest (Dixit and Pyndick, 1994) and its role for credit constraints, is rather more plausible than the model of firm behaviour used in the New Keynesian structure. But progress is always going to be incremental.

---

[6] Which it could do by representing a wider range of financial assets.

Conclusions

New models offer interesting and new ways of looking at the world. I certainly hope that the comparisons I have made between some of the results generated by disaggregate models and those which would arise from an analysis based on a representative agent should warn us against putting too much weight on findings from the latter. It seems to me unlikely that a model based on a representative agent can deal satisfactorily with the effects of credit constraints or fears about availability of credit. There is a sharp contrast between the conclusions from aggregate and disaggregate modelling about desirable tax structures. Looking at some issues I have covered, analysis of myopia indicates, if nothing else, that at least one of the assumptions of "behavioural economics" may not have a major impact on saving behaviour. And study of education indicates to me that costs, returns and uncertainty cannot account for everything; here, at least, other factors appear to be important.

I also think that development of models of this type is going to help us understand the nature of the shocks which have affected the international economy over the last five years. Shocks in the banking sector have proven to be self-reinforcing, perhaps through the multiplier process described by Gai, Haldane and Kapadia (2011). Tightening of credit and fear of tight credit has affected both firms and households and, in the United Kingdom at least, this may be limiting the scope for recovery. Disaggregate modelling offers a coherent means of exploring this latter issue.

But, let me end with a warning. I would not want you to infer that I believe the development of the "true" model lies just around the corner. For that, you should look at the end of the rainbow, where there is also a pot of gold for whoever has built it.

References

Alan, S., T. Crossley and H. Low. (2012). "Saving on a Rainy Day; Borrowing for a Rainy Day". University of Cambridge. Mimeo.

Bank of England (2012). "The Distributional Effects of Asset Purchases". *Bank of England Quarterly Bulletin.* Pp. 254-266.

Bean, C. (2012). "Pension Funds and Quantitative Easing". Speech given to National Association of Pension Funds. 23[rd] May.

Bloom, N. (2009). "The Impact of Uncertainty Shocks". *Econometrica.* Vol 77. Pp 623-685.

Chamley, C. (1986). "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives". *Econometrica.* Vol 54. Pp 607-622.

Conesa, J.C., S. Kitao and D. Kreuger (2009). "Taxing Capital? Not a bad Idea after all". *American Economic Review.* Vol 99. Pp. 25-48.

Deaton, A. (1991). "Saving and Liquidity Constraints". *Econometrica.* Vol 59. Pp. 1221-48.

Diamond, P. (2008). "Behavioural Economics". *Journal of Public Economics.* Vol 92. Pp 1858-1862.

Diamond, P. and B. Köszegi. (2003). "Quasi-hyperbolic Discounting and Retirement". *Journal of Public Economics.* Vol 87. Pp. 1839-1972.

Dixit, R.K and R.S. Pindyck. (1994). *Investment under Uncertainty.* Princeton University Press.

Fang, H. and D. Silverman. (2009). "Time Inconsistency and Welfare Programme Participation: Evidence from the NLSY". *International Economic Review.* Vol 50. Pp 1043-1077.

Gai, P. A. Haldane and S. Kapadia. (2011). "Complexity , Concentration and Contagion". *Journal of Monetary Economics. .* Vol 58. Pp 453-470.

Goodhart, C. (2007). "Whatever became of the Monetary Aggregates". *National Institute Economic Review.* No. 200. Pp. 56-61.

Gornemann, N, K. Kuester and M. Nalajima (2012). "Monetary Policy with Heterogeneous Agents". Federal Reserve Bank of Philadelphia. Presented to NBER Summer Institute.

Laibson, D., A. Repetto and J. Tobacman. (2007). "Estimating Discount Functions from Lifecycle Consumption Choices". Department of Economics. Harvard University.

Lucas, R.E. (1990). "Supply-side Economics: an Analytical Review". *Oxford Economic Papers.* Vol 42. Pp 293-316.

McIntosh, S. (2006). "Further Analysis of the Returns to Academic and Vocational Qualifications". *Oxford Bulletin of Economics and Statistics.* Vol 68. Pp. 225-251.

Miles, D. (2012). "Asset Prices, Saving and the Wider Effects of Economic Policy". Speech given to pro- Manchester Business Conference. 1[st] March.

Mirrlees, J.A. (1971). "An Exploration in the Theory of Optimum Income Taxation". *Review of Economic Studies.* Vol 38. Pp. 175-208.

Oh, H. And R. Reiss. (2011). "Target Transfers and the Fiscal Response to the Great Recession". NBER Working Paper No 16675.

Plosser, C.I. (2012). "Macro Model and and Monetary Policy Analysis".
Bundesbank-Federal Reserve Bank of Philadelphia Spring Research Conference.
http://www.philadelphiafed.org/publications/speeches/plosser/

Sefton J, and J. van de Ven (2009), "Optimal Design of Means Tested Retirement Benefits*,
Economic Journal, Vol:119, Pages:F461-F481

Van de Ven, J. (2010). "The Effects of Myopia on Retirement Savings Decisions",
National Institute of Economic and Social Research.

Van de Ven, J. and M.R. Weale. (2012). "The Decision to take up Adult Education: a Dynamic Programming
Analysis". Presented to LLAKES Conference. October

Woodford (2011). Simple Analytics of the Government Expenditure Multiplier.
American Economic Journal.: *Macroeconomics.* Pp 1-35.