Utility Maximising Design of Means Tested Retirement Benefits*

James Sefton† and Justin van de Ven‡

Abstract

The design of welfare benefits is a tricky business. In this respect, James Meade placed particular emphasis on the importance of avoiding excessive distortions to the price of labour. Nevertheless, Meade noted that means testing is likely to be desirable in view of the "hideously expensive" cost of universal benefits provision – he conjectured that a 50% withdrawal rate on welfare benefits might be appropriate. In this study we take a fresh look at the role of means testing in the provision of retirement benefits in the United Kingdom. We use an articulated rational agent model of the household to explore the effects on expected lifetime utility of alternative budgetary neutral pensions arrangements. In this context, we find that extensive means testing of retirement benefits is preferred, consistent with the conjectures stated by Meade. Our analysis highlights the importance of taking into account the distortions associated with alternative methods of benefits financing when considering pensions reform.

JEL Classifications: C61, D91, I38, J26

Key Words: optimal taxation, social insurance, means testing, pensions reform.

1 Introduction

> *Every citizen, rich or poor, receives the same tax-free Basic Income but the Surcharge on the first slice of other income is the equivalent to a withdrawal of part of the Basic Income, not pound for pound but, say, one pound for every two pounds earned. The Surcharge may weaken but will not eliminate the incentive to earn more income; it can however, enormously reduce the cost of the Basic Income.* "Can we learn a 'third way' from the Agathotopians?", Meade (1993, p. 94).

Means testing has become something of a vilified concept in the contemporary debate regarding the design of retirement benefits in the United Kingdom. That this debate should focus primarily on the distortions associated with benefits withdrawal is understandable, given the perceived need to encourage higher savings in the context of an aging population. It does, however, risk over-stating the case against means testing. As noted by Meade (1993, p. 94), the provision of a universal welfare benefit can be "hideously expensive", and means testing is a tool for limiting that fiscal burden.

*Presented at the conference in honour of James Meade, Bank of England, 12-13 July, 2007.

†The Business School, Imperial College London, and NIESR. jsefton@niesr.ac.uk

‡NIESR. jvandeven@niesr.ac.uk

The role of means testing in the provision of retirement benefits was one of the central issues addressed by the Pensions Commission in its second report, Pensions Commission (2005). Based on the premise that means testing discourages private saving, the recommendations made by the Pensions Commission – which have now been largely accepted by UK Government in subsequent pensions White Papers – will halt the spread of means tested retirement benefits in the near future, and reduce their absolute importance in the longer term. In this paper, we take a fresh look at the question of means testing in the reform of state retirement benefits. We assume that household preferences can be described by the rational agent life-cycle model, and explore the welfare effects of alternative retirement benefits design. Our analysis suggests that a representative household would prefer a system in which retirement benefits are subject to extensive means testing. Self-evidently this finding is conditional on our analytical context and assumed model of preferences. In these respects, the policies that we consider for analysis are framed following the example given by the Pensions Commission, and our welfarist approach is standard in the economic literature that considers the intertemporal allocation of resources (Browning & Crossley (2001)). To the extent that the emerging consensus regarding pensions policy reform in the UK is coalescing about the Pensions Commission's recommendations, our results would consequently appear to warrant careful consideration. They suggest that means testing cannot be rejected by a simple appeal to its effect on incentives to save; a decision must be made taking into account the distortions associated with alternative methods of benefit financing.

The contemporary policy debate has raised a number of important criticisms of means testing, and it is impractical to make adequate allowance for all of these here. Before moving on to discuss the specifics of our study, it consequently seems appropriate to address the question: *what is our analysis missing?* One of the most frequently cited criticisms of means testing in general is that the associated benefits suffer from poor rates of take-up, potentially due to a social stigma attached to their receipt. Means testing is also most naturally adapted for application to benefits that are provided at the family or household level, which runs counter to the political trend toward welfare provision to the individual. As summarised by Atkinson (1996, p. 68), "In short, means-testing is economically inefficient, provides an incomplete safety net, and takes social policy backwards rather than forwards." Furthermore, it has been argued that means testing encourages tax-evasion by setting "the interest and the duty of individuals at variance" (Slemrod (2007)). We should be clear that the current study focuses upon the distortionary effects of benefits means testing, and of benefits financing more generally. No allowance is made to reflect concerns regarding benefits take-up, the intra household allocation problem, or issues associated with tax implementation, and to the extent that these are important, our analysis is likely to overstate the case in favour of benefits means testing.

Our analysis is based upon an articulated and carefully calibrated rational agent model of the household. This model assumes that decisions regarding consumption and labour supply are made to maximise expected lifetime utility, subject to various practical constraints, including age specific liquidity limits, the imposed tax and benefits system, and uncertainty regarding labour incomes and time of death. Alternative specifications for the tax and benefits system are compared on the basis of their respective implications for expected lifetime utility. The importance that is attached to uncertainty in our analysis relates the current study to the optimal tax literature that is concerned with the *social insurance-efficiency* trade-off. This may be contrasted with the more extensive optimal tax literature that explores the *equity-efficiency* trade-off, following the seminal work by Mirrlees (1971). In the current context, taxation serves the role of hedging individuals against some form of uninsurable risk. As noted by Varian (1980, p. 51), "The motive for redistribution here is not a desire for equity *per se*, but rather a desire for *social insurance*" (emphasis in the original). Focussing upon the social insurance-efficiency trade-off has generally been found to result in marginal rates of taxation higher than those reported in the equity-efficiency literature, and to provide a motive for a progressive rate structure.[1] This last conclusion of the social insurance literature is easy to understand, in view of the fact that extreme values of income are taken by the social insurance literature to be the product of very good or bad luck, against which a risk averse individual would like to insure.

Observations drawn from the optimal tax literature that is concerned with the social insurance-efficiency trade-off suggest that the optimal progressivity of the tax structure is increasing in risk aversion and uncertainty, and decreasing in prudence.[2] The first two of these effects are intuitive enough, and are likely to play an important role in the results that we report here. Regarding the negative relation between prudence and preferences for progressivity, we note that both Varian (1980) and Low & Maldoom (2004) consider analytical contexts in which precautionary behavioural responses tend to augment the tax base, and Low & Maldoom (2004) cite this as the mechanism behind the negative relation that they identify. In the context of intertemporal consumption and labour supply decisions, however, it is not clear that this relationship between prudence and the tax base will continue to hold. Attanasio et al. (2005), for example, report that precautionary behaviour tends to motivate increased labour supply and saving early during the working lifetime, and earlier retirement later in life.

[1] Mirrlees (1971) finds that the optimal tax schedule is approximately linear, and much of the subsequent literature has focussed upon how the optimal schedule is influenced by alternative analytical assumptions. A stylised observation of this literature is that it is difficult to motivate a progressive marginal tax rate structure. That marginal tax rates of the optimal tax schedule will not be everywhere increasing in the absence of uncertainty follows immediately from the following often-cited observation. When individual heterogeneity arises purely from differences in (unobservable) ability, then any tax schedule that imposes a non-zero marginal tax rate on the most able individual is Pareto dominated by one that imposes a lower marginal tax rate since this will make the most able person better off, and the remainder of the population made no worse off. See, Seade (1977).

[2] See, for example, Varian (1980) and Low & Maldoom (2004). Risk aversion is defined in the usual way ($-u''/u'$). Preferences are described as exhibiting "prudence" if the sign on the third derivative is positive ($u''' > 0$), which implies decreasing risk aversion with consumption. In the context of additively time separable utility, Leland (1968) proves that individuals will only undertake precautionary saving if their preferences exhibit prudence.

If the tax base is more sensitive to variations in labour supply later in life, then the results of Attanasio et al. (2005) suggest that prudence may actually reduce the tax base, with consequences for the welfare maximising design of taxation.[3] Although we do not explore this issue further here, it is an effect that warrants further research.

We explore the welfare effects of varying three policy parameters – the value of means tested retirement benefits; the withdrawal rate applied to means tested retirement benefits; and a generic tax rate. Following the Pensions Commission[4], the analysis is conducted holding the remainder of the tax and benefits system fixed, as it was applied at a single point in time. As is usual in the optimal tax literature, we restrict our attention to those policies that have the same budgetary effect in aggregate, so that our policy analysis is subject to two degrees of freedom. It is of note here, that budget neutrality is imposed from a generational accounting perspective (see Auerbach et al. (1994)), so that each cohort is subject to the same aggregate (present discounted) tax burden, measured over its entire life-span.[5] Sensitivity to the methods of benefits financing is conducted, with tax adjustments modelled as either a poll-tax, a proportional tax (a coarse attempt at raising the revenue in a distributionally neutral manner), or an increase in the basic rate of tax. We show that our principal finding – that some means testing of retirement benefits is preferred on an expected lifetime utility basis – is not sensitive to the choice of financing rule. However the severity of the means test, and particularly the size of the means tested benefit is. This second conclusion of the analysis highlights an important, but often neglected point: decisions regarding pensions policy reform should take full account of how the required revenue is raised. Reforming the pensions system in isolation could lead to *third* best policy decisions.

The policies that are considered for analysis and associated motivation in terms of the contemporary pensions debate are described in Section 2. The intuition regarding the welfare effects of means tested retirement benefits is developed using a two period model in Section 3. Section 4 provides a brief description of the dynamic programming model that is used to undertake the analysis, and Section 5 reports numerical results. A summary and directions for further research are presented in the conclusion.

2 Policies Considered for Analysis

The policy alternatives that are considered for analysis were selected to reflect the contemporary debate regarding the design of retirement benefits in the UK. This section begins with a brief review of that

[3] This relation would be further exaggerated if the elasticity of labour supply is higher toward the end of the working life. This is likely to be the case for the model used by the current study, which assumes that labour incomes are influenced by an experience effect. This experience effect tends to provide a stronger motive to supply labour early in the simulated lifetime, and was included to match the model against profiles for labour supply estimated from survey data. See Section 4 for further discussion. See, also Mirrlees (2006), pp. 15-16, for discussion of further complications arising due to the intertemporal nature of the problem.

[4] See, for example, Appendix C of the Pension Commission's first report, and Appendix F of their second report.

[5] Hence, we do not take into account the transitional effects of budgetary reform, or the fiscal implications of the changing age structure of the population. These issues were addressed at length by the Pensions Commission.

debate, before providing a detailed description of the policy alternatives considered for analysis.

2.1 The contemporary pensions debate

Debate regarding the appropriate design for the future of pensions provision in the UK has coalesced around the recommendations made by the Pensions Commission in its second report. The Pensions Commission was established to identify a long-term strategy for pensions policy, following a period of extensive reform that suggested considerable uncertainty regarding a fundamental policy question: is it best to target benefits at poorer pensioners through means testing, or encourage (and possibly compel) everyone to save sufficiently for their own retirement? In short, the contemporary debate can be described as focussing upon the balance between social and private insurance. As a way of illustrating the degree of uncertainty, we shall briefly contrast the various initiatives.

In July 1998, the government announced an above inflation increase in (means tested) income support for pensioners as a first step to establishing a Minimum Income Guarantee (MIG). This income support was designed to 'top up' income in retirement to the level defined by the MIG, and was consequently subject to a 100% withdrawal rate on private income (a loss of £1 of benefit for every £1 of private income until the benefit was exhausted).[6] Every subsequent budget saw the level of the MIG increase relative to the basic pension until, by April 2003, the MIG exceeded the basic pension by £24.65 per week for single person (or 5.6% of average gross full-time employment income) – up from £5.75 per week in 1998. Furthermore, as the level of the basic pension was set to rise in line with prices, and the Government declared that the MIG would rise in line with wages, the importance of means testing was set to increase into the future. Over this period, opposition to means testing of retirement benefits gathered momentum, primarily focussing upon the argument that means testing discouraged individuals to save for their retirement.[7] Partly as a response to this opposition, the government introduced the Pension Credit from October 2003, which reduced the withdrawal rate on means tested retirement benefits in excess of the basic pension from 100% to 40%.

A parallel set of pensions policy reforms have their roots in the contributory principle: that pension benefits are the return to earlier savings or contributions, and not an unconditional right paid for out of general tax revenue. The pensions system in the UK was founded on the contributory principle when it was first introduced almost one hundred years ago. However, successive amendments to the system have seen the basic state pension – the central pillar of the UK pensions system – became more similar

[6] The measure of 'income' that is referred to here includes imputations for the income generated by assets, calculated separately in practice through the imposition of assets tests.

[7] For example, Mary Francis Director General of the Association of British Insurers stated 'Our concern is that as means-testing is extended ... too many people will conclude that it doesn't pay to save' (March 2003, ABI Press Release); Frank Field, MP and Minister for Welfare Reform 1997-8, 'the Government's ... emphasis on means tests now seems misplaced. It is quite clear that the message has gone out ... that savings can damage their retirement income' (Hansard, 4th June 2003); Institute for Public Policy Research published the Report, 'New Contract for Retirement' that stated 'there is clear cut evidence that a policy framework that relies heavily on means-testing retirement benefits is flawed'.

to a universal benefit than a truly contributory entitlement. To augment the contributory nature of the state pension system, the State Earnings Related Pension Scheme (SERPS) was introduced from 1978. SERPS was replaced by the State Second Pension (S2P) in April 2002, which weakened the contributory principle underlying SERPS by providing more generous terms to individuals on low incomes.

In 2002 the Government established the Pensions Commission to review of the UK pensions system and make recommendations. The Commission's proposals have two principal targets. Firstly, they establish a form of 'participation income', Atkinson (1996), whereby the basic pension and the S2P together are made flat rate, and their (aggregate) value increased relative to thresholds associated with means tested benefits.[8] With time, these reforms are designed to reduce the number of people dependent on means tested benefits. Secondly, and incidental to this study, the reforms introduced a centralised system of low cost personal savings accounts into which individuals will be automatically enrolled and contribute 8% of their salary (though they can choose to opt out).

2.2 A stylised model of the policy debate

The terms of reference of the Pension Commission were to 'review the regime for UK private pensions and long-term savings' and to make recommendations 'on whether there is a case for moving beyond the current voluntarist approach', Department for Work and Pensions (2002). These terms of reference were interpreted broadly. The commission made recommendations about the reform of both the public and private pension systems. However the mandate required that the recommendations be made in the context of the rest of the UK tax and benefit system. In this paper we recognise this constraint, and look at the design of a stylised social security system whilst assuming that the incentives embedded in the rest of the tax and benefit system remain broadly fixed. However any increase in generosity of a pension system must be paid for, and so it is impossible to leave the tax system entirely unchanged in a closed structural model. We therefore look at the design of the UK public pensions system under three alternative tax raising scenarios; a poll tax, a uniform increase in all marginal rates, and an increase in only the basic tax rate. We begin by describing the stylised pensions system that is considered for analysis before discussing the three tax raising scenarios that are used to balance the fiscal budget.

Benefits during retirement

It is useful to begin by defining terms, and providing details regarding the rates and thresholds assumed under our default policy environment. In 2003 the basic pension was £77.45 per week for an individual and £ 123.80 per week for a couple. For comparison, the median earnings of all full-time employees

[8] We use the term 'participation income' here to denote rights to receipt that are allocated on the basis of a set of individual circumstances that extends beyond simple labour market activity.

in the UK in the winter of 2003/04 was £369 per week.[9] We assume that all households receive a full basic state pension from age 65, conditional on the number of adults in the household. This assumption reflects two off-setting considerations: on the one hand not all households are eligible for the full basic pension from age 65, and on the other hand we assume that all households are contracted out or pay National Insurance tax at the lower rebate rate.[10]

In addition to the basic pension, households may receive some income support. Under the default policy environment, this is specified to reflect the the Minimum Income Guarantee payable in 2003, equal to £102.10 per week for an individual and £155.80 per week for a couple. This additional income is subject to a 100% withdrawal rate. In view of the fact that all households are considered to receive the full basic pension, means tested benefits are received by individuals under the base scenario up to the point that their private incomes exceed £24.65 per week for a single, and £32.00 per week for a couple.

We stylise the debate about pension policy in the UK, as a debate about both the maximum potential value of the means tested benefit, and the rate of benefits withdrawal. Thus the early years of the Labour Government – which came to power in 1997 – were all about increasing the maximum potential value of the MIG. The introduction of Pension Credit, by contrast, was a change in the associated withdrawal rate. Consistent with the Pension Credit reform, we only consider changes in the withdrawal rate for means tested benefits in excess of the basic state pension. We stylise the pension commission reform as a reduction in the withdrawal rate to zero. It should be noted, however, that we do not consider increases in the value of the basic state pension, in addition to the rates and thresholds of means tested benefits. Hence, the analysis would need to be extended to take the full policy reforms enacted by the 2006 pensions white papers into account.

The government budget constraint

The various tax and benefits systems that are considered for the retired are not budgetary neutral when taken in isolation. How the government chooses to treat the fiscal effects of pensions reform can have an important welfare implications. In common with the optimal tax literature, we assume that taxes must rise in order to meet any fiscal deficit due to the provision of a retirement benefit.

Our analysis is based upon numerical simulations derived from our model of the UK household sector, which is described in Section 4. In this model, each household is considered to make consumption and labour supply decisions at annual intervals to maximise its expected life-time utility, given uncertainty about both its future income and lifespan. We assume that the economy is made up of large number of

[9] Reported in Table 34 of the *Labour Force Survey (LFS) Historical Quarterly Supplement*, published by the Office for National Statistics.

[10] See Sefton et al. (2006) for further discussion.

these households, all of whom take the economic environment as given.

We consider a population that is sufficiently large so that there is no uncertainty at the level of population aggregates. We also make a small economy assumption (appropriate for the UK), so that interest rates and wage rates do not respond to changes in aggregate labour supply and savings. These assumptions help to simplify the task of imposing a government budget constraint.

To avoid dynamic effects associated with inter cohort redistribution, the government budget constraint is applied in terms of the generational account of the youngest cohort. The generational account is defined as the net present value of all transfers made between all members of this cohort and the government during the cohort's life. Thus the generational account is the present value of all taxes paid minus the present value of all benefits received, so that our welfare analysis takes into consideration the full costs of alternative retirement benefit arrangements.

There are clearly a number of alternative tax adjustments that could be considered for analysis. We consider three, each of which is applied exclusively to households under the age of 65. We refer to the first as a "poll tax" – a vertical translation of the tax and benefit schedule that takes the same fixed amount in every period of a household's working life (ages 20-64), regardless of any temporal fluctuations in gross income. Thus, if $£x$ of tax were paid on an income of $£y$ in the base economic environment, then $£(x+d)$ would be paid following the application of a poll tax of $£d$, *for all* $£y$. Gradient methods were adopted to search for the appropriate value of d to adopt for any given specification of the retirement benefits system. This tax adjustment has the advantage that it does not distort relative prices. It is, however, a regressive method of balancing the budget.

The second approach to balancing the budget is referred to as a "proportional tax". This adjustment amounts to a shift in marginal tax rates, resulting in a flattening (or a steepening) of the tax schedule in every period of a household's working life. We note, however, that to omit the possibility of marginal tax rates in excess of 100%, the proportional tax adjustment was not applied to the lowest incomes that are subject to means tested benefits during the working lifetime. Specifically, tax schedules are simulated as four piece linear splines by the model, and the proportional tax adjustment was only applied to the upper three segments of each of these splines. We chose this as our second case as it is approximately distributionally neutral (in terms of relative incomes), and is in our view politically feasible. It does, however, affect relative prices during the working lifetime.

Our third and final scenario, was neither non-distortionary or distributionally neutral. It was chosen with an appeal to pragmatism. One of the most commonly adjusted parameters on the UK tax system is the basic rate of tax.[11] This has been adjusted regularly over the last 30 years and has varied between

[11] Income taxes in the UK are imposed by a standard multi-step function, with four marginal rates. No tax is applied to income under an individual's "personal allowance". Three rates of marginal tax are then applied, increasing with gross income, from a "lower rate", to a "basic rate", and finally a "higher rate".

33% and 20%. Our third scenario therefore adjusts the basic rate of tax to bring the generational account into balance.

3 Implications from a Two Period Model

The first period of the two period model corresponds to the working life of a household, and the second period to retirement. Let w_1 denote the total (disposable) resources available to a household at the beginning of period 1 (full-time labour income plus any physical capital endowment). Wealth at the beginning of retirement, w_2, is then first period savings plus the return to investment:

$$w_2 = (1 + r)(w_1 - e_1) \geq 0$$

where r denotes the real (after tax) rate of return to saving, and e_1 is total household expenditure in period one.

In retirement, all households are assumed to receive the universal basic pension, p. In addition there is a means tested pension benefit, p_c, which is withdrawn at the (constant) marginal rate t_m with regard to wealth, w_2. Consumption in the second period is consequently equal to:

$$c_2 = p + w_2 + \max(p_c - t_m w_2, 0)$$

Assume that intertemporal preferences are time separable, and that intratemporal preferences are smooth and concave in both consumption and leisure in period one, and in consumption in period two. Furthermore, assume that intratemporal preferences between consumption and leisure in period one are homothetic and that both choice variables are selected from a continuous, closed and bounded set. Then the intertemporal utility maximisation problem can be solved by two-stage budgeting, where consumption and leisure in the first period are linear functions of total expenditure (for internal solutions).

A great deal is revealed about the behavioural and welfare effects of means tested retirement benefits by this simple analytical example. Suppose that the two period utility function, specified in terms of total expenditure in periods 1 and 2, is defined by:

$$u = \frac{1}{1 - \frac{1}{\gamma}} \left(e_1^{1-1/\gamma} + \delta e_2^{1-1/\gamma} \right) \tag{1}$$

where δ is the discount rate and γ the intertemporal elasticity of total expenditure. Given this model structure, and parameter values that reflect those of the fully articulated model considered for the remainder of the paper, Figure 1 reports the effects on incentives, welfare, and the value of retirement benefits received, of a fall in the pensions withdrawal rate from 70 to 40 percent (an arbitrary policy adjustment considered for the purposes of exposition).

Panel A of Figure 1 distinguishes between four income/wealth bands, based upon the resources held in period one. The panel indicates that decreasing the withdrawal rate on means tested retirement

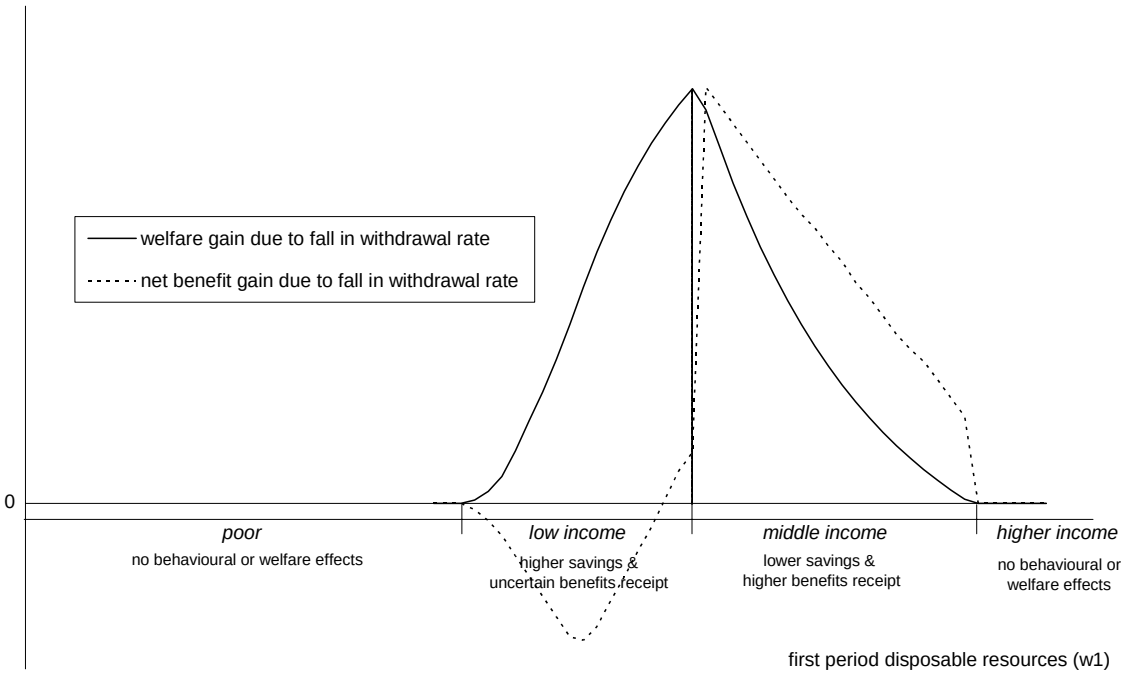

Panel A: welfare and budget effects by first period wealth endowment

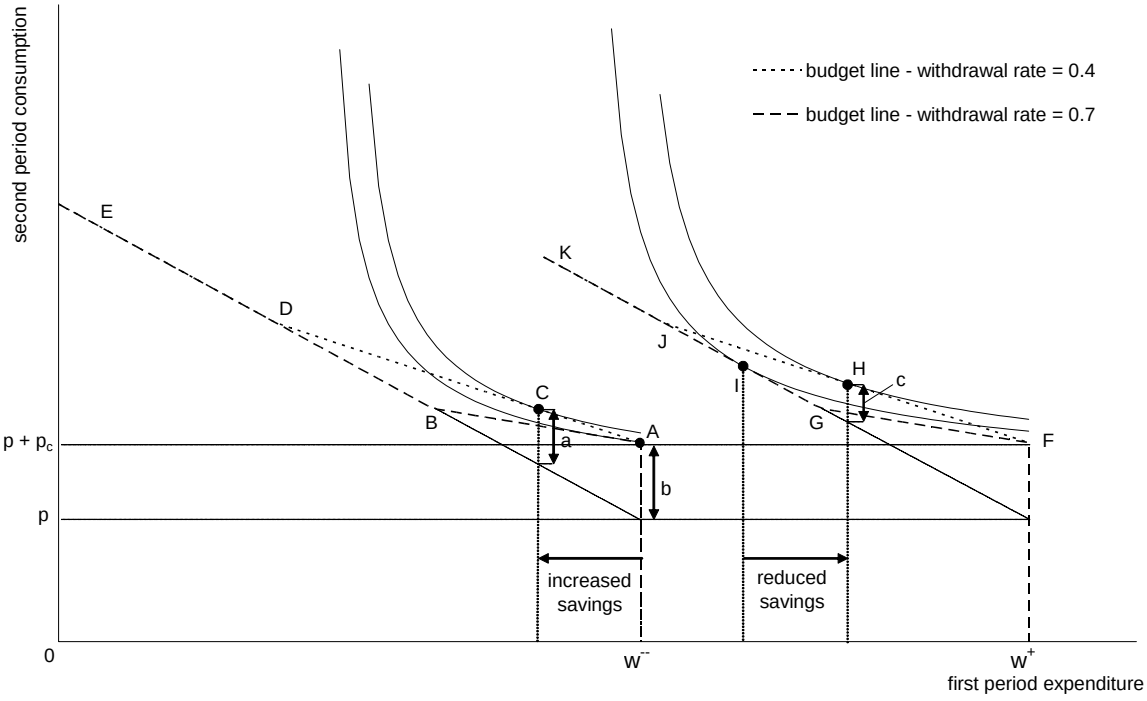

Panel B: savings responses of low and middle income agents

Notes: Author's calculations based upon similar specification of preferences as considered for the fully articulated model described in Section 4
The authors may be contacted for full details

Figure 1: Behavioural, Welfare, and Budgetary Effects of a Decrease in the Withdrawal Rate on Means Tested Retirement Benefits from 70 to 40 percent

benefits has no effect on the welfare or benefits received by a household at either extreme of the wealth distribution. At the bottom of the distribution, a household is unaffected by the considered policy change because the pension benefits are sufficiently generous (relative to their available resources during the working lifetime) to dissuade any private provision for retirement, even when the withdrawal rate is low. In contrast, at the top of the income distribution, the public insurance programme appears relatively insubstantial, so that the households concerned make their own provisions for retirement regardless of the imposed withdrawal rate. In between these two extremes, the welfare gain associated with the fall in benefit withdrawal rates increases with first period wealth to a maximum before falling back away. This profile arises due to the alternative behavioural responses of a household at different points in the distribution – labelled low and middle income in the figure – with the distinguishing threshold between the two defined where the welfare gain as a result of the pension reform is highest.

Panel A of Figure 1 clearly has important implications for the aggregate effects of means testing on both the government budget and expected lifetime utility. Before we discuss these further, it is worth taking a moment to discuss the determinants of the profiles that are reported in Panel A of the figure.

First, note that the income effect of a fall in the withdrawal rate on means tested pensions tends to decrease first period savings, and that this is off-set by the associated price effect. Given the analytical specification described above, it can be shown that a marginal decrease in the withdrawal rate on means tested pensions will decrease savings at a sufficiently high w_1 relative to $(p + p_c)$, and increase savings otherwise.[12] This is because the income effect increases with the amount that a household chooses to save when subject to the higher pensions withdrawal rate.

Suppose that means tested pensions are specified such that the price effect associated with a decrease in the pension withdrawal rate dominates for any agent with an internal solution to the utility maximisation problem that implies receipt of means tested benefits prior to the change. This is indicated in Panel B of Figure 1 by the household with first period wealth endowment $w_1 = w^-$. Then a utility maximising agent will decrease their savings in response to a marginal decrease in the pension withdrawal rate only if they are made better off following the policy change by moving onto means tested benefits having not received benefits previously. This circumstance is indicated in Panel B by the household with first period wealth endowment $w_1 = w^+$. The agent that is indifferent between receiving and not receiving means tested benefits when subject to the higher withdrawal rate defines the transition between those who increase and those who decrease their savings in response to the policy change.

It can also be shown that – when the price effect dominates as described above – the welfare gain

[12] Given the model as described here, it can be shown that, $\frac{d\hat{w}_2}{dt_m} = \frac{(1-\gamma)\delta^\gamma(1-t_m)^\gamma(1+r)^\gamma}{\left[(1-t_m)+\delta^\gamma(1-t_m)^\gamma(1+r)^{\gamma-1}\right]^2} w_1 - \frac{1+\gamma\delta^\gamma(1-t_m)^{\gamma-1}(1+r)^{\gamma-1}}{\left[(1-t_m)+\delta^\gamma(1-t_m)^\gamma(1+r)^{\gamma-1}\right]^2}(p+p_c)$ where $\hat{w}_2$ denotes utility maximising first period saving.

associated with a marginal decrease in the pension withdrawal rate is increasing in first period wealth for any agent with an internal solution that implies receipt of means tested benefits prior to the change.[13] This is reflected in Panel A of Figure 1 by the upward sloping section of the graph denoted for low income agents. In contrast, middle income agents are those who did not receive means tested benefits prior to the fall in the withdrawal rate, but do receive benefits following the policy change. In their case, the additional utility derived from the policy change falls with w_1, because utility maximising savings are an increasing function of w_1, which limits the receipt of means tested benefits at any given withdrawal rate. This explains the downward sloping section of the graph displayed for middle income agents in Panel A.

What do the profiles reported in Panel A of Figure 1 reveal regarding the drivers that underly the welfare effects of means tested retirement benefits with which the current study is concerned? In the introduction, care is taken to note that the current study focuses upon the implications for expected lifetime utility of alternative retirement benefits arrangements, considered in the context of labour income uncertainty: we are concerned with the social insurance-efficiency trade-off of pensions policy design. With regard to the two period model that is presented here, this can be interpreted as uncertainty over the first period wealth endowment, w_1.

The welfare profile in Panel A of the figure reveals that the influence on expected lifetime utility of a fall in the withdrawal rate of means tested retirement benefits depends crucially upon the probability that is assigned to being in different parts of the endowment distribution. In the example considered here, a high probability of being in the "low" or "middle" segments of the distribution would imply a large welfare gain due to the fall considered for the benefit withdrawal rate. Similarly, the impact on the aggregate government budget also depends upon how the population is assumed to be distributed, relative to the benefits in question. As described in Section 2, our numerical analysis assumes that tax adjustments during the working lifetime are made to off-set the budgetary impact of alternative pensions policy counterfactuals, which also influence welfare. Hence, distributional issues are central to any evaluation of means tested benefits policy, and the need to combine both the immediate incidence effects of policy with the more general problem of benefits financing exaggerates the complexity of associated welfare analysis.

[13] For agents in receipt of means tested benefits and some savings, a marginal decrease in the benefits withdrawal rate will necessarily increase lifetime utility. The way in which this increase varies with first period wealth is described by the second order condition: $\frac{d^2V}{dt_m dw_1} = -\frac{1}{\gamma}\hat{e}_1^{-\frac{1}{\gamma}-1}\frac{d\hat{e}_1}{dt_m}$, where $V = U(\hat{e}_1, \hat{e}_2)$ denotes utility at the optimum. Since $\hat{e}_1, \gamma \geq 0$ by assumption, if the price effect dominates (so that $d\hat{e}_1/dt_m > 0$), then $d^2V/(dt_m dw_1) < 0$. Hence, the rise in utility associated with a marginal decrease in the pensions withdrawal rate is larger as w_1 increases.

4 An Articulated Rational Agent Model of the Household

No closed form solution exists for the problem that is considered in Section 5, and so analysis is based upon simulations derived from a dynamic programming model of the household.[14] The model considers lifetime consumption and labour supply decisions of households in annual increments, from age 20 to the maximum potential age of 110. Full details of the model are reported in Sefton et al. (2007), where we focus upon the behavioural responses to the introduction of the Pension Credit (in contrast to the welfare effects of pensions design, which is the current subject of concern). In that study we show that calibration of the model against survey data ties the model parameters down fairly tightly. This section consequently provides a brief review of the main features of the model, and the interested reader is referred to that paper for further details.

4.1 The utility function

Expected lifetime utility of household i at age t is described by the time separable function:

$$U_{i,t} = \frac{1}{(1-1/\gamma)} E_t \sum_{j=t}^{110} u\left(\frac{c_{i,j}}{\theta_{i,j}}, l_{i,j}\right)^{1-1/\gamma} \delta^{j-t} \phi_{j-t,t} \tag{2a}$$

$$u\left(\frac{c_{i,j}}{\theta_{i,j}}, l_{i,t}\right) = \left(\left(\frac{c_{i,j}}{\theta_{i,j}}\right)^{(1-1/\varepsilon)} + \alpha^{1/\varepsilon} l_{i,t}^{(1-1/\varepsilon)}\right)^{\frac{1}{1-1/\varepsilon}} \tag{2b}$$

where E_t is the expectations operator, $\phi_{j-t,t}$ is the probability of living to age j, given survival to age t, and $\delta = 0.97$ is the discount factor, which is assumed to be the same for all households and time independent. Household i is considered to choose consumption at each age t, $c_{i,t}$, and leisure $l_{i,t} \in \{l_{FT} = 0.7, l_{PT} = 0.85, 1\}$ to maximise equation (2a), subject to their wage, wealth, and the relations that are considered to govern the intertemporal dynamics of their environment. From age 65 (the current state pensionable age of men in the UK), all households are forced to retire if they have not already done so, in which case $l_{i,t} = 1$ for all i an $t \geq t_{SPA} = 65$. $\theta_{i,t} \in R^+$ is adult equivalent size based upon the McClements' scale[15], which is considered to evolve deterministically.

$\gamma = 0.2$ is the intertemporal elasticity of substitution (of total expenditure), and $\varepsilon = 0.58$ is the (period specific) elasticity of substitution between equivalised consumption $c_{i,t}/\theta_{i,t}$ and $l_{i,t}$. The constant $\alpha = 1.63$ is referred to as the utility price of leisure. The partial differential of equation (2a) with respect to consumption, $c_{i,t}$, and leisure, $l_{i,t}$, is given by:

$$U_{cl,t} = \left(\frac{1}{\varepsilon} - \frac{1}{\gamma}\right) \frac{u_{c,t} u_{l,t}}{u_t^{1+1/\gamma}} \delta^{j-t} \phi_{j-t,t} \tag{3}$$

[14] For examples of studies of retirement behaviour based on structural DP models see, for example, Gustman & Steinmeier (1986), Rust & Phelan (1997), and French (2005).

[15] See McClements (1977) on the McClements' equivalence scale, and Balcer & Sadka (1986) and Muellbauer & van de Ven (2004) on the use of this form of adjustment for household size in the utility function.

where we use the standard notation to denote partial derivatives. As within period utility and the marginal utilities of consumption and leisure are all positive, consumption and leisure are direct substitutes for the analysis that is presented here. Furthermore, the model parameters considered for analysis are associated with an average intertemporal elasticity of substitution for consumption of 0.376 which lies within the range of values considered by the literature.[16]

4.2 The wealth constraint

Equation (2a) is considered to be maximised, subject to a wealth constraint, $w_{i,t} \geq 0$. We define total wealth as:

$$w_{i,t} = \begin{cases} w_{i,t-1} + \tau\left(l_{i,t-1}, r_{i,t-1}w_{i,t-1} + x_{i,t-1}, n^a_{i,t-1}, n^c_{i,t-1}, t-1\right) - c_{i,t-1} & \text{if } t \lesseqgtr t_{SPA} \\ (1-\eta)\left[w_{i,t-1} + \tau\left(l_{i,t-1}, r_{i,t-1}w_{i,t-1} + x_{i,t-1}, n^a_{i,t-1}, n^c_{i,t-1}, t-1\right) - c_{i,t-1}\right] & \text{if } t = t_{SPA} \end{cases} \tag{4}$$

where $r = 4\%$ is the real interest rate, $x_{i,t}$ is private non-property income, and $\tau(.)$ is the tax and benefit function. $n^a_{i,t}$ and $n^c_{i,t}$ refer respectively to the numbers of adults and children in a household, which – like the equivalence scale – are considered to evolve deterministically with age. We abstract from the asset allocation problem and define $w_{i,t}$ as total household wealth. In the first period of life (age 20), all households are considered to have zero wealth, $w_{i,20} = 0$ for all i.

At age $t = t_{SPA}$, a proportion, $\eta = 0.5$, of household wealth is annuitised at an actuarially fair rate χ.[17] During the working lifetime, $t < t_{SPA}$, $x_{i,t}$ defines household labour income, equal to $\varphi\left(l_{i,t}\right) h_{i,t}$, where $\varphi\left(l_{i,t}\right)$ is the proportion of the full-time employment wage, $h_{i,t}$, earned: $\varphi\left(l_{FT}\right) = 1$; $\varphi\left(l_{PT}\right) = 0.28$; and $\varphi\left(1\right) = 0$.[18] $h_{i,t}$ is considered to evolve following the stochastic process that is described in subsection 4.3. During retirement, $x_{i,t}$ is equal to the annuity income generated by private pensions:

$$x_{i,t} = \begin{cases} \varphi\left(l_{i,t}\right) h_{i,t} & \text{if } t < t_{SPA} \\ \eta\chi\left[w_{i,t-1} + \tau\left(l_{i,t-1}, rw_{i,t-1} + x_{i,t-1}, n^a_{i,t-1}, n^c_{i,t-1}, t-1\right) - c_{i,t-1}\right] & \text{if } t = t_{SPA} \\ x_{i,t-1}\frac{\left(0.5+0.5\left(n^a_{i,t}-1\right)\right)}{\left(0.5+0.5\left(n^a_{i,t-1}-1\right)\right)} & \text{if } t > t_{SPA} \end{cases}$$

The function τ is a stylised representation of the UK tax and benefit system, and is described in subsection 4.4.

4.3 Income dynamics

In the first period of the simulated lifetime, age 20, each household is allocated a wage, $h_{i,20}$, via a random draw from a log-normal distribution, $\log(h_{i,20}) \sim N\left(\mu_{20}, \sigma^2_{20}\right)$, where $\mu_{20} = 5.443$ and

[16] Calculated at population averages for consumption (£453 per week), leisure (0.740), and the equivalence scale (1.837 * 300) between ages 25 and 60, weighting each age equally.

[17] η is important in the current context because the proportion of wealth that is not annuitised is considered to be exempt from means testing. Of the various holdings that are included in the composite asset, $w_{i,t_{spa}}$, two classes are omitted from the eligibility (income and wealth) tests that have been applied by means tested retirement benefits: owner occupied housing, and the first £6,000 of additional wealth. In the current context, we assume that these exempt assets account for 50% of $w_{i,t_{spa}}$.

[18] These differences between part-time and full-time employment are specified to reflect averages reported in survey data by the Office for National Statistics. See, *Labour Force Survey (LFS) Historical Quarterly Supplement*, Table 38.

$\sigma_{20} = 0.4$. Thereafter, wages are generated using the stochastic process described by the equation:

$$\log h_{i,t} = \beta \log h_{i,t-1} + \kappa \frac{(1 - l_{i,t-1})}{(1 - l_W)} + f(t-1) + \omega_{i,t} \tag{5}$$

where $f(t)$ is an age-dependent wage growth term, $\beta = 0.99$ accounts for time persistence in earnings, $\omega_{i,t} \sim N\left(0, \sigma_\omega^2\right)$, $\sigma_\omega = 0.18$ is a household specific disturbance term, and $\kappa = 0.05$ is the return to another year of experience. The parameters assumed for the wage generating process imply strong intertemporal persistence, and a large annual wage premium in return for employment, which reflects the difficulties that are commonly experienced in capturing labour participation rates at young ages.

4.4 The tax function

The function τ is a stylised representation of the UK tax and benefit system, described as a (continuous) four piece linear spline function of the household's pre-tax income, that is its property income $r_i w_{i,t}$ plus non-property income $x_{i,t}$, its equivalent size $n_{i,t}^a$ and $n_{i,t}^c$, and its age, t. The age dependency assumed for the tax function divides the lifetime into three periods: the working lifetime $t < t_{IB} = 55$, early retirement $t_{IB} \leq t < t_{SPA} = 65$, and retirement $t_{SPA} \leq t$. During the working lifetime, the base specification of the tax function – prior to any adjustment made for budget balance – is specified to reflect profiles reported in the April 2003 edition of the *Tax Benefit Model Tables* (TBMT) issued by the Department for Work and Pensions.[19] The profiles considered take into consideration the impact of income taxes, National Insurance Contributions, the Child Benefit, the Working Tax Credit and the Child Tax Credit. Although this list omits a great deal of the detail of the UK tax and benefits system, it does include the principal schemes that affected healthy families with children during 2003.

The simulated tax function for ages $t_{IB} \leq t < t_{spa}$ depends upon private income, employment status, age, and demographic composition. Simulated households that choose to supply labour for any t, $t_{IB} \leq t < t_{spa}$, are treated in the same way as during the working lifetime (described above). The tax treatment applied to a simulated household that chooses not to supply labour and is aged $t_{IB} \leq t < t_{MIG} = 60$, is specified to reflect the Incapacity Benefit and income taxes as they stood in 2003/4; between ages $t_{MIG} \leq t < t_{spa}$ the tax function is specified to reflect the Pension Guarantee (identical for the alternative policy counterfactuals considered here) and income taxes. Tax treatment throughout ages $t_{IB} \leq t < t_{spa}$ is also subject to the tax adjustments made to ensure budget balance.

The tax function during retirement, $\tau(.)$, $t \geq t_{spa}$, is discussed at length in subsection 2.2. It is specified to reflect the effect of income taxes, the basic state pension, and means tested benefits. Each of the policy counterfactuals considered for analysis hold the rates and thresholds of income taxes and the basic pension fixed, as they were applied in 2003/04. Alternative policy counterfactuals differ during retirement in the value of, and withdrawal rates applied to means tested benefits.

[19] See http://www.dwp.gov.uk/asd/tbmt.asp.

5 Results

The study is framed in terms of the following thought experiment. Suppose that individuals from a birth cohort are asked at age 19 – before they have learned anything about their individual specific circumstances – to select their preferred policy specification from a menu that is limited to revenue neutral alternatives. What properties is the preferred policy option likely to possess?

In undertaking the current analysis we have chosen to trade analytical tractability for empirical detail. The dynamic programming model described in Section 4 was consequently used to derive measures of expected lifetime utility following application of the incentive compatibility constraint, for each policy counterfactual, over every feasible combination of private income and wealth at age 20. We assume that, prior to the beginning of the simulated life, all individuals are identical. They are then randomly endowed with a wage and wealth pair at age 20, after which they are considered to behave as predicted by our simulation model. Expected utility prior to entry into the model is consequently the same for all individuals, so that the policy decision that we are interested in here is taken as behind a veil of ignorance. We calculate the expected utility at age 19 for each considered policy combination by taking the expected lifetime utility at age 20 for any combination of wealth and wage that is calculated numerically by the model, and combining this with the joint distribution of wealth and wage at age 20 derived from survey data. A policy combination, A, that generates a higher expected utility at age 19 than some other policy combination, B, is then considered to be preferred to policy B.

The analytical model that we adopt for the current analysis trades analytical tractability for descriptive detail. The disadvantage is that we are unable to describe the analytical properties of our problem – we do not obtain *necessary* let alone *sufficient* conditions for a global optimum. Consequently, the numerical results that we report can at best be described as local optima, subject to the methods used in their calculation. Thankfully, it appears that our analytical problem is well behaved, in the sense that the objective function over which we search for a maximum does not exhibit substantial discontinuities, nor does it appear to possess multiple peaks. This is an advantage, as the problem is numerically demanding, even in our highly stylised context.

5.1 Numerical Results

We now assess the welfare effects of alternative retirement benefits arrangements, and the associated budgetary implications. We should note that the incentive compatibility constraint underlies the current analysis, so that the simulations upon which the statistics reported here are based fully incorporate the savings and labour supply responses to means tested retirement benefits that are alluded to in section 3.

Table 1: Preferred Pension System of Representative Household as a Function of the Approach used to Balance the Government Budget Constraint

tax adjustment to balance generational accnt	means tested benefit value	benefit withdrawal rate
Poll tax of £107.40 per year (2003 prices)	1.8 x value in 2003	70%
Proportional tax increase of 2.7%	3.9 x value in 2003	70%
Increase in the basic tax rate of 2.85%	3.0 x value in 2003	60%

Given the assumptions that are embodied by our analytical model, we find that expected lifetime utility obtains a local maximum under the combination of policy parameters reported in Table 1. To illustrate these results we summarise the case when taxes are raised proportionally to balance the government budget constraint. We find that the pension system that maximises the utility of our representative household raises the value of means tested benefits by 390%; thus for a single pensioner this is equivalent to increasing the maximum award in 2003 from £24.65 a week to £96.14 per week, and for a couple of raising the award from £32 per week to £124.80 per week. (Note that this award is received in addition to a basic state pension of £77.45 per week for a single and £123.80 per week for a couple in 2003). These increased means-tested benefits are, however, withdrawn at a rate of 70%; thus 70p of benefits is lost for every £1 of private pension income. Further to pay for this pension benefit, it is necessary to raise the marginal tax rates across the entire income distribution by 2.7%. The results for the two other approaches to balancing the government budget constraint can be understood in an identical manner.

We can immediately make two deductions from these results, which we shall explore in more detail later. Firstly these results, perhaps surprising given the actual recommendations of the UK Pensions Commission, suggest a significant role for means testing in the provision of benefits to the retired. Rather than decreasing the dependence of the pensions system on means tested benefits, we find that our representative household would prefer to increase them. Secondly, the actual size of the increase depends critically on the approach adopted to raise the revenue for these increased benefits. If the extra revenue is raised by a poll-tax, an approach that can be understood as either introducing no additional redistribution or social-insurance into the tax system, then the desired increase in the level of pension benefits is small. If the extra revenue is raised by a proportional increase in all tax rates, an approach that can be understood as being more redistributional or introducing more social-insurance into the tax system, then the desired increase in the level of pension benefits is large. The third case, a rise in the basic rate of tax, lies between these two extremes. As we find the proportional increase in all taxes generates a higher increase expected lifetime utility for our representative household than the

other two approaches, these results are all consistent with our representative household desiring more redistribution or social insurance than is currently implicit within the UK tax system *considered over the entire lifetime*.

We now investigate the results obtained for the proportional tax increase in more detail, noting that the other two cases considered for analysis are qualitatively the same. In Figure 2, we report the sensitivity of welfare effects of alternative pension arrangements, measured in terms of compensating variations relative to the system as it stood in 2003. Panel A of the figure reports contour maps for expected lifetime utility at age 19, relative to the rate of benefits withdrawal, and the value of means tested benefits as a multiple of their value in 2003. Panels B and C complement Panel A by reporting, respectively, horizontal and vertical cross-sections of the surface reported in Panel A, taken through the local maximum. Panels B and C also plot average welfare effects across the cohort of households at age 65. Unsurprisingly, the welfare effects on retired households are monotonic in the generosity of the benefits system.

The first observation is that there is a clear local maximum for the welfare profile measured at age 19; in fact the profile of welfare effects appears to be concave over the class of pension systems considered for analysis, for values between 0 and 5 times contemporary levels and withdrawal rates between 0 and 100%. The statistics reported in Figure 2 do, however, reveal that the measured welfare effects are reasonably similar for withdrawal rates between 50 and 90% about the identified local maximum, and between 3.5 and 4.1 times their current levels.

The absolute value of compensating variations are much larger when measured from age 65, than when measured from age 19, underscoring the fact that the elderly are likely to express a stronger preference over retirement benefits than the young. This is in part driven by the fact that the effect of any change in pensions policy on younger households is dampened by the discount rate that is assumed to govern intertemporal preferences, and in part because older households are not subject to the full fiscal burden of any reform to retirement benefits. It is of note that the compensating variation measured at the identified local maximum from age 65 is 4.6 times the annual increase in benefits to couples and 5.9 times the annual increase for singles, which are both less than half the value of 13.7 that is associated with an actuarially fair annuity in the analysis. This is because few households receive the entire increase in the means tested benefit that is identified for the local maximum, due to the high withdrawal rate that is imposed. For households from age 19, the value of the compensating variations reported in Figure 2 reveal that the alternatives considered have a small impact on expected lifetime utility. These observations emphasise the difficulties that may arise when basing pensions policy reform on voter preferences.

The effects of the government budget on proportional tax adjustments for alternative specifications

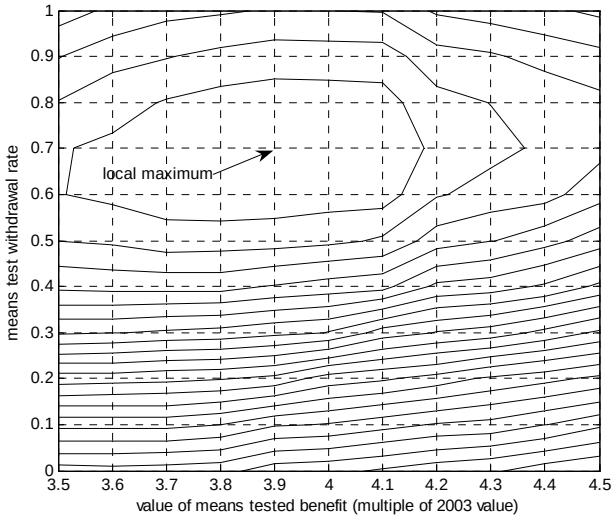

Panel A: varying both means test benefit value and withdrawal rate

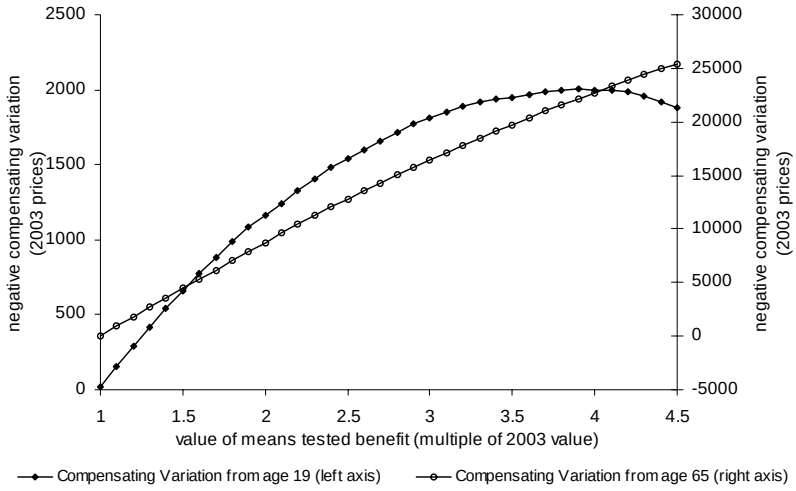

Panel B: holding means test withdrawal rate fixed at 70%

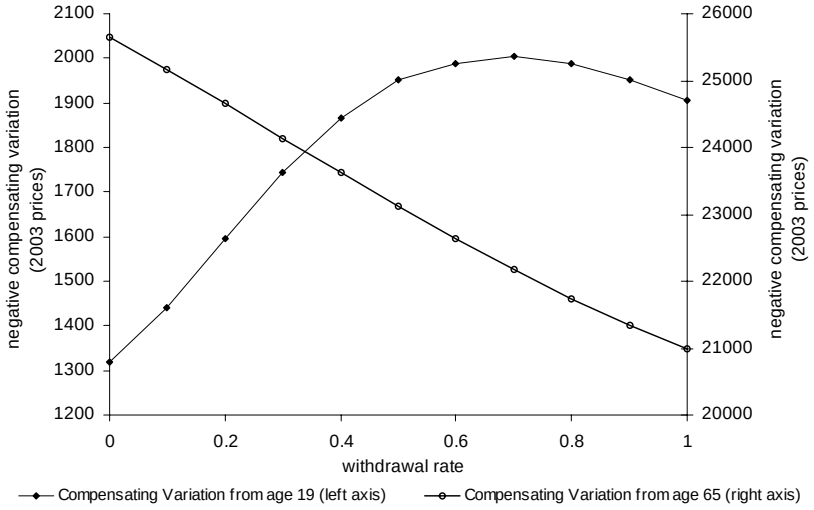

Panel C: holding the value of means tested benefits fixed at 3.9 x 2003 value

Figure 2: Effects on Expected Lifetime Utility when Budget Neutrality is Maintained by Proportional Tax Adjustments

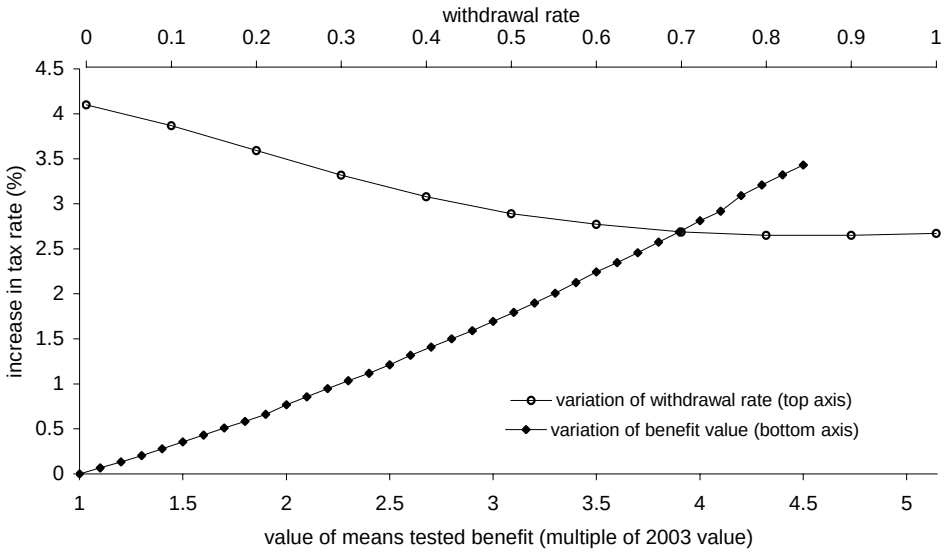

Notes: series denoted "variation of withdrawal rate" holds benefit value equal to 3.9 times 2003 value
series denoted "variation of benefit value" holds withdrawal rate equal to 70%

Figure 3: Proportional Tax Rate Adjustments to obtain Budget Neutrality between Alternative Policy Counterfactuals

of the retirement benefits system are reported in Figure 3. Again, the adjustments that are reported are cross-sections taken through the local maximum identified for welfare, as in Panels B and C of Figure 2. These tax adjustments underly the welfare effects that are discussed above; and underlying the tax adjustments are behavioural responses to the tax and benefits system. Hence, Figure 4 reports population averages for behavioural responses of labour supply and savings. We discuss the behavioural responses and their influence on the tax adjustments required for budget balance, before relating the sum to the welfare effects that are reported above.

The two panels of Figure 4 reveal that, whereas varying the value of means tested benefits tends to have a similar impact on behaviour throughout the range reported here, the impact of varying the rate of benefit withdrawal is highly non-linear. This is attributable to the fact that when increasing the value of means tested benefits, wealth effects for those in receipt of the benefit dominate, leading to reduced savings and labour supply in aggregate. In contrast, changing the rate of benefit withdrawal while holding the value of the benefit fixed exaggerates the role of price effects, which tend to have opposing influences on households at different points in the income / wealth distribution as discussed in section 3. Panel B of Figure 4 indicates that these militating behavioural responses are dominated by higher income households with regard to saving, and by lower income households with regard to labour supply – reducing benefits withdrawal rates encourages increased labour supply from those on means tested benefits prior to the policy change, are reduced saving by new benefits recipients after the policy change.

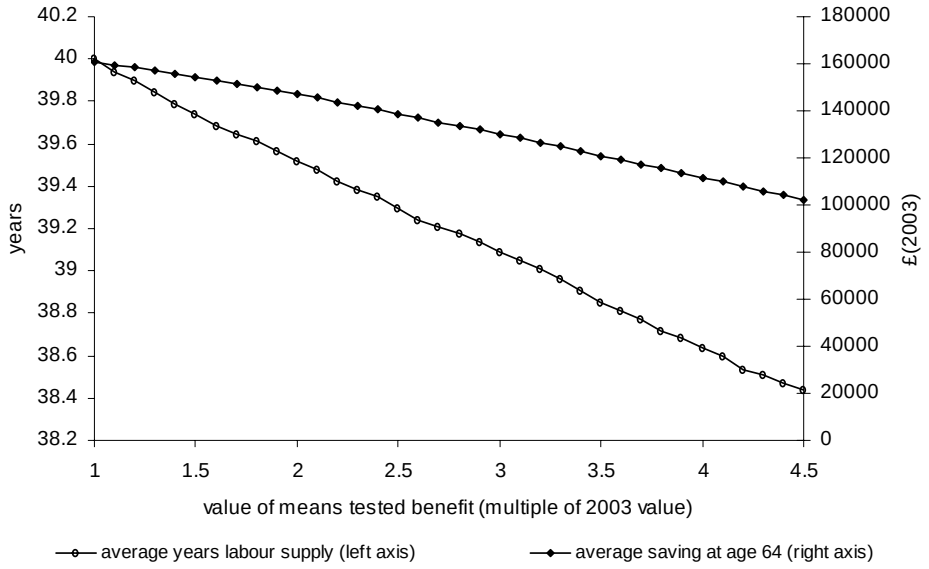

Panel A: holding means test withdrawal rate fixed at 70%

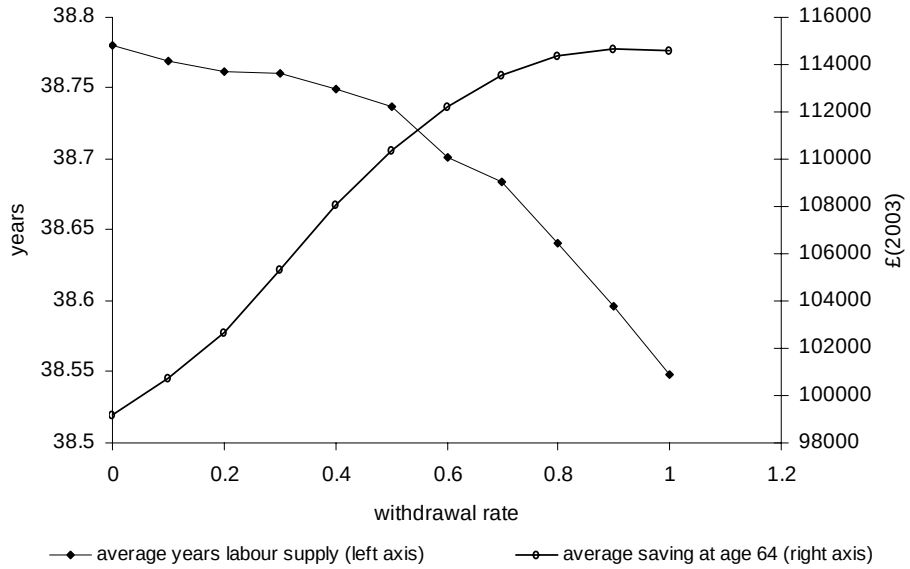

Panel B: holding the value of means tested benefits fixed at 3.9 x 2003 value

Figure 4: Aggregate Behavioural Responses to Pensions Policy Counterfactuals

These behavioural responses explain the budget balancing tax changes that are reported in Figure 3. The dominance of wealth effects when the value of means tested benefits is increased gives rise to a monotonic increase in the proportional tax adjustment required to maintain budget balance. It is of note, however, that the monotonic increase of the proportional tax is accelerating with respect to the value of the means tested benefit. This is because of the coincident falls in labour supply and saving, which tend to reduce the effective tax base. Indeed, this effect is a principal factor in limiting the size of the preferred means tested benefit. In contrast, the tax adjustment for budget balance is not monotonic in the withdrawal rate applied to means tested pensions – tax rates fall when withdrawal rates are reduced from 100 to 90 percent, and rise again as the withdrawal rate is reduced below that level. This is because, at very high rates of benefits withdrawal the behavioural responses of those on means tested benefits dominate, resulting in an increase in aggregate tax revenue when the withdrawal rate is marginally reduced. It is this factor that limits the preferred withdrawal rate on means tested benefits.

6 Conclusions

In relation to the design of a redistributive tax and benefit system, Meade (1993) suggested that particular care should be exercised to avoid distortions to the price of labour. To avoid distortions to labour, Meade noted, redistribution must be undertaken with respect to non-labour income. Three approaches are suggested: a Property Owning Democracy (reduction of existing wealth inequality through effective redistribution); the Social Ownership of Property (the state takes possession of national assets by drawing down government debt); a Welfare State (income from pay converted to income unrelated to pay through the imposition of taxes and benefits). In relation to the third of these alternatives, Meade (1993, p. 94) had the following to say:

> Another form of tax employed by the Agathotopians to finance a really adequate Basic Income is the imposition of a special surcharge on the first slice of every citizen's income other than the Basic Income itself. This is a form of levy which in effect withdraws part of the Basic Income as a citizen's other income increases, a levy which is additional to the ordinary current rate of tax on income. This turned the Agathotopian Basic Income into something which is half way between a fully conditional social benefit and a fully unconditional social benefit. With a fully conditional benefit a citizen without other income is given the full-scale benefit, but this benefit is reduced pound for pound as the citizen's other income increases. This is the cheapest way to ensure that everyone has a minimum guaranteed income. Benefit is paid only so far as it is needed to bring income up to the

> minimum level. But it has the well-known effects of removing any incentive to earn any additional income, so long as any such income will be matched by an equivalent reduction in social benefit. On the other hand a completely unconditional social benefit removes this disincentive effect, since the same tax-free social benefit is received regardless of the level of other income. But it is hideously expensive in that it hands out free of tax an adequate social benefit to every citizen however rich or poor. The Agathotopian Basic Income with a Surcharge on the first slice of other income falls between these two extremes.

In contrast to this view, one of the most important findings of the optimal tax literature (Atkinson & Stiglitz (1976)) is that, in the absence of uncertainty, if household utility is separable in goods and leisure, then there is no welfare motive for the application of indirect taxation. In the current context, this finding is important because, as Atkinson and Stiglitz note, a means test on retirement benefits can be interpreted as an intertemporal expenditure tax. Consistent with this last view, contemporary debate regarding the future of pensions policy in the UK is now resolving about the recommendations made by the Pensions Commission, which will limit the role played by means testing in the provision of benefits to the retired.

In this paper we take a fresh look at the debate in the United Kingdom regarding the role of means testing in the provision of retirement benefits. Our analysis is based upon an articulated rational agent model in which the lifetime is divided into annual increments from age 20 to 110, and households are considered to choose their labour supply and consumption at each age as though to maximise their expected lifetime utility, given their individual specific circumstances and the uncertainty that describes the intertemporal evolution of labour incomes and the time of death. Beginning with a stylised description of the UK tax and benefits system as it stood in 2003, we search over a class of redistributive systems that differ in the value means tested retirement benefits, the rate of the benefits withdrawal, and a series of alternative tax adjustments that ensure budget neutrality. We compare the alternative tax and benefit counterfactuals with reference to their respective effects on the expected lifetime utility at the beginning of life, as implied by our rational agent model. On the basis of this analysis, we find that the preferred policy counterfactual from our considered alternatives implies extensive means testing of retirement benefits – expected lifetime utility is found to achieve a local maximum with a means tested benefit value of between 1.8 and 3.9 times the that was applied in 2003, and a withdrawal rate on private income of between 60 and 70%, depending upon the tax adjustment that is considered to ensure budget neutrality. These results are in stark contrast to the recommendations made by the Pensions Commission, and relate more closely to the views expressed by Meade as quoted above.

We do not, however, wish to give the impression that the results reported here represent an appropriate basis for reforming UK pensions policy. Our results describe the logical implications of a set of highly restrictive assumptions, and it is not at all clear how these results would change in the context of alternative – and more reasonable – assumptions. One of the most important limitations of the current study, for example, is the limited policy space over which we have explored. Although it is well beyond the scope of this conference paper to search over the entire intertemporal domain for the optimum tax and benefits policy (if such an optimum exists at all), the sensitivity of the preferred pensions policy counterfactuals that we have identified here to the tax adjustments that are considered for budget neutrality highlight the difficulties in extrapolating to policy relevant conclusions. Our results could appropriately be interpreted as implying that, given our assumed preferences and the uncertainty that we have considered to influence intertemporal circumstances, a more progressive tax structure, *as considered from a lifetime perspective*, would be preferred. In relation to means testing of retirement benefits, our results are sufficiently striking to suggest that further evaluation of the long-term direction of UK pensions policy is in order.

It would be of obvious interest to relax many of the policy parameters that we have held fixed for the current analysis: the basic state pension, and the structure of taxes and benefits during the working lifetime are two prominent examples. More fundamentally, it would be interesting to consider whether there remains a role for means testing of retirement benefits if labour income taxation was allowed to adjust to maximise expected lifetime utility when intratemporal preferences between consumption and leisure are not separable, and in the context of labour income uncertainty. This latter study would shed light on how far the optimal tax result of Atkinson and Stiglitz (1996) may be carried over into a more complex – and realistic – environment. These issues remain for further research.

References

Atkinson, A. B. (1996), 'The case for a participation income', *The Political Quarterly* **67**, 67–70.

Atkinson, A. B. & Stiglitz, J. E. (1976), 'The design of tax structure: direct versus indirect taxation', *Journal of Public Economics* **6**, 55–75.

Attanasio, O., Low, H. & Sanchez-Marcos, V. (2005), 'Female labor supply as insurance against idiosyncratic risk', *Journal of the European Economic Association* **3**, 755–764.

Auerbach, A. J., Gokhale, J. & Kotlikoff, L. J. (1994), 'Generational accounting: a meaningful way to evaluate fiscal policy', *Journal of Economic Perspectives* **8**, 73–94.

Balcer, Y. & Sadka, E. (1986), 'Equivalence scales, horizontal equity and optimal taxation under utilitarianism', *Journal of Public Economics* **29**, 79–97.

Browning, M. & Crossley, T. F. (2001), 'The life-cycle model of consumption and saving', *Journal of Economic Perspectives* **15**, 3–22.

Department for Work and Pensions (2002), Simplicity, Security and Choice: working and saving for retirement. DWP White Paper, Command Number 5677.

French, E. (2005), 'The effects of health, wealth, and wages on labour supply and retirement behaviour', *Review of Economic Studies* **72**, 395–427.

Gustman, A. L. & Steinmeier, T. L. (1986), 'A structural retirement model', *Econometrica* **54**, 555–584.

Leland, H. E. (1968), 'Saving and uncertainty: the precautionary demand for saving', *Quarterly Journal of Economics* **82**, 465–473.

Low, H. & Maldoom, D. (2004), 'Optimal taxation, prudence and risk-sharing', *Journal of Public Economics* **88**, 443–464.

McClements, L. (1977), 'Equivalence scales for children', *Journal of Public Economics* **8**, 191–210.

Meade, J. E. (1993), *Liberty, Equality and Efficiency*, New York University Press, New York.

Mirrlees, J. (2006), *Welfare, Incentives, and Taxation*, Oxford University Press, Oxford.

Mirrlees, J. A. (1971), 'An exploration into the theory of optimum income taxation', *Review of Economic Studies* **38**, 175–208.

Muellbauer, J. & van de Ven, J. (2004), Estimating Equivalence Scales for Tax and Benefits Systems. NIESR Working Paper, 229.

Pensions Commission (2005), *A New Pension Settlement for teh Twenty-First Century – the second report of the Pensions Commission*, The Stationery Office, Norwich.

Rust, J. & Phelan, C. (1997), 'How social security and medicare affect retirement behavior in a world of incomplete markets', *Econometrica* **65**, 781–831.

Seade, J. (1977), 'On the shape of optimal tax schedules', *Journal of Public Economics* **7**, 203–235.

Sefton, J., van de Ven, J. & Weale, M. (2006), Means Testing Retirement Benefits: fostering equity or discouraging savings? NIESR Discussion Paper 283.

Sefton, J., van de Ven, J. & Weale, M. (2007), Means testing retirement benefits: fostering equity or discouraging savings? forthcoming in the Economic Journal.

Slemrod, J. (2007), 'Cheating ourselves: the economics of tax evasion', *Journal of Economic Perspectives* **21**, 25–48.

Varian, H. R. (1980), 'Redistributive taxation as social insurance', *Journal of Public Economics* **14**, 49–68.