

10. FRAMING MATTERS: NON-STANDARD FISCAL POLICY

10.1 Introduction

In the introductory chapter we encountered the general demand function $x(p,m,f)$, where p represented the price vector, m was income and f was the frame or context. I use the term non-standard fiscal policy to denote the manipulation of the f term in order to achieve government goals. For instance, Cullis and Jones, 2000, suggest that because individuals who pay for an item are more likely to consume it, then prescription charges may help individuals complete a course of treatment. This chapter is devoted to exploring the potential uses of this kind of non-standard policy in more detail. Some illustrations come from the theory of ‘libertarian paternalism’ proposed by Camerer et al, 2002. Much of the focus is on savings and pensions policies. For that reason I also include discussion of some relevant intertemporal decision-making anomalies.

No presumption is made about the goals of government: they might be to maximize welfare or to pursue some more limited objective (e.g. ‘raise savings’) which potentially has a welfarist justification, or they might arise from the aims of rent seeking groups etc. There is also no assumption that these policies can be used to manipulate individual behaviour at will. Policies based on framing must have their limits, just as standard fiscal policies do.

The sorts of instruments which might be tried are shown in Table 10.1. In the first column we have the class of anomaly or heuristic; in the second we have the policy lever and then the final column provides an example. So, for instance, given the tendency for individuals to stick with the status quo, setting the default as opt-in for pensions may lead to much higher savings levels compared to the case where individuals have to choose to opt in. As Camerer et al, 2002, note such policies have the merit that they do not penalise the rational agent, since typically the actual cost of switching is low. On the other hand they produce potentially large benefits for the individual who has difficulty making the commitment to save or become insured. Meanwhile, in his discussion of mental accounting, Thaler, 1981, points out that individuals are much more sensitive to out-of-pocket expenses than they are to the opportunity costs. As a result, individuals may overuse cars relative to public transport, because the depreciation cost of using the former is not salient. Conversely, individuals may undervalue education, if the opportunity costs of not attending college are hard to appreciate.

To explore further aspects of Table 10.1, there are four major elements of this chapter: first a discussion of labelling policies uses the the UK’s Winter Fuel Allowance as a motivating example; second we examine inconsistency in intertemporal decision-making which is then discussed further in the context of savings policy. The final substantive section is more speculative and focuses on framing effects in tax policy.

Table 10 1. Non-Hicksian Policies

Anomaly	Policy	Example
Reference Points and Anchoring	Manipulating the endowment point Using the choice set to manipulate reference points.	Use of defaults in insurance and pension legislation. Providing dominated options
Mental Accounting	Offering commitment devices to reduce temptation Labelling to create accounts Using the non-equivalence of out-of-pocket and opportunity costs. Exploiting the non-sunk nature of costs	Penalties for early withdrawal from savings accounts Hypothecated taxes and expenditures. Educational attendance subsidies; information on depreciation costs of car use Prescription charges to encourage consumption of full course of treatment.
Information Biases	Exploiting the availability and representativeness heuristics	'Tokenism' in visual advertising, restrictions on celebrity advertising

10.2 Creating mental accounts.

As intimated, mental accounting is a potential motive for a variety of policies that when viewed through the lens of standard theory defy easy explanation. One of these is the U.K.'s Winter Fuel Allowance, about which some background is in order.

Along with Ireland, the UK has the highest percentage of excess deaths during the winter months in Europe. Mortality rates regularly rise during periods of cold weather and deaths are concentrated in the over 65s. The exact link between temperature and excess death rates for the UK are unclear. Hypothermia accounts for a relatively small number of the deaths, most of which are attributed to respiratory or circulatory diseases (Aylin et al, 2001). Hypothesised and investigated causes of the deaths include poor quality housing stock, lack of central heating, pensioner behaviour and activity rates as well as poverty.

The Winter Fuel payment has been presented as a means of reducing excess mortality by allowing pensioners to spend more on their heating bills. Introduced as a token payment in 1997 and extended significantly four years later, the allowance is currently a lump-sum payment of £200 payable annually to those households claiming it, with an extra £100 for households where at least one person is aged 80 or over. Households who are not already in receipt of state benefits have to claim, but there is no income contingency and no requirement that the money be spent on fuel bills. Once a claim is paid in one year, the payment is automatically made in subsequent years unless the Department for Work and Pensions (DWP) is alerted to a material change in circumstances. Originally the allowance was available to men over 65 and women over 60 living in the UK, but following a ruling by the European Court of Human Rights, all individuals over 60 can claim. The UK government does not produce figures on the numbers potentially eligible. However, currently the total number of households receiving benefit is over 8.5 million.¹

10.2.1 Theory.

Within standard Hicksian theory the allowance is equivalent to a lump sum grant contingent solely on age (and the fact that the payment is claimed). As a result, the impact on heating should be minimal. As a guide, note that (own) estimates derived from the FES for 2001/2 suggest that the income elasticity of fuel expenditure is rather small at around 0.28. For an average two-adult retired household in the second lowest decile of the income distribution this translates into £2.78 per year from a £200 rise in income or about five pence per week. For households in the top fifth of the income distribution, the rise is only £1.50 per year.

More fundamentally the existence of the Minimum Income Guarantee (now the Pension credit) meant that when the Winter Fuel Allowance was intro-

¹From 2003-04, the Food and Expenditure Survey (FES) - a nationally representative survey of UK households, asks questions about receipt of the benefit. From that it appears that above the age of 62, over 95% of eligible households are in receipt. Slightly lower uptake levels are recorded for households where the oldest person is 61. See Munro, 2006a.

duced there was pre-existing vehicle for giving pensioners a lump-sum rise in income. The use of this vehicle would have avoided the significant costs of running the Winter Fuel programme. In other words, as instrument of policy, the allowance was dominated by adjustments to pre-existing programmes.

If the programme is hard to explain on welfarist or fuel poverty reduction grounds, it is equally difficult to produce a sensible political economy explanation that does not at some point invoke bounded rationality. The three main groups of actors relevant to the programme are pensioners, non-pensioners and elected politicians. Suppose there was a political outcome in which Δ was paid in cash to all pensioners at a net cost of $\Delta + c$ to non-pensioners. Compare this to an outcome in which pensioners receive Δ at a cost $\Delta + c'$ where $c' < c$. Non-pensioners would prefer this second outcome provided their preferences satisfy non-satiation and provided the burden for each individual was lower under the second arrangement. Meanwhile pensioners would be indifferent or prefer the second arrangement if they were at all altruistic. Given this, politicians would also prefer the second arrangement, provided that their payoff was increasing in the utility of citizens.

One possible though not necessarily plausible explanation is based on asymmetric information along the lines of Coate and Morris, 1995. For instance, non-pensioner citizens might be against a lump-sum subsidy to pensioners but in favour of a targeted heating subsidy.² If they are not able to distinguish between a heating subsidy proper and the Winter Fuel payment then there may exist a political equilibrium in which the Winter Fuel payment is made. One problem with this argument is the high level of publicity given to the programme: the government advertises it heavily, individual Labour party members of parliament produce localised figures on its uptake in their election material and newspapers carry reminders to pensioners to claim. More fundamentally, in a world where some actors can choose whether information is asymmetric or not, asymmetric information cannot be a part of any equilibrium if that is not compatible with incentives. In other words, if non-pensioners were unsure of the true nature of the payment, then there would be clear gains from “political arbitrage” in the sense that there would be returns to an opposition politician or indeed newspapers from passing on the information. Yet the Winter Fuel payment remains popular, largely unchallenged in the U.K.

An alternative theory is that there is no citizens are actually bounded in their rationality, but that many voters believe that others are so. This is entirely possible, but then we are into the world of bounded rationality.

The most sensible explanation from within a bounded rationality framework is that the Winter Fuel payment creates a mental account out of which the income elasticity of expenditure is higher than in the standard model. The discussion of mental accounting in Chapter Two ended with the observation that there was as yet no good formal model of it, although prospect theory could go some way towards explaining many of the features of mental accounts. So I begin there. We shall see though that there are problems with using the riskless version of prospect theory.³

²In the final section of this chapter I argue that labels on taxes and subsidies may affect behaviour. One might therefore reasonably argue that a targeted heating subsidy can have a mental accounting effect along with the usual incentives to consume provided by a lower consumption price.

³I am grateful to Henrik Kleven for alerting me this point.




Figure 10-1 Prospect Theory and the Winter Fuel Payment.

Let us assume the following.

1. Individuals assess the utility of different options using Kahneman and Tversky style valuation functions. I shall take it that all valuation functions are twice differentiable, except at 0.
2. In particular $w(x)$ is a concave function representing the utility gain from x (monetary) units of heating expenditure. Gains in income are valued by the function $u(x)$ while $v(x)$ is used to value losses in income or expenditures. These functions satisfy the usual properties of valuation functions in prospect theory. In particular we take $u(0) = v(0) = 0$ for all x , and $u(x) < v(x)$ and $u'(x) < v'(x)$, where $'$ indicates a first derivative.
3. Where cash grants, y , are not named or associated with heating, then heating expenditures and cash grants are not integrated into one mental account. As a result x is chosen to maximize: $w(x) + v(-x) + u(y)$.
4. Where cash grants are named or otherwise associated with heating, then heating expenditures and cash grants are integrated into a single mental account. As a result, x is chosen to maximize, V where :

$$V = w(x) + v(-x + y) \quad \text{for } x > y$$

$$V = w(x) + u(y - x) \quad \text{for } y > x.$$

The basic model is illustrated in Figure 10-1 for the specific case where $y > x$. It can be seen that $u(y) + v(-x) < u(y-x)$.

Where expenditure and grant are not integrated, the level of x is set by the condition.

$$w'(x) = v'(x) \tag{10.1}$$

Note that this is independent of y which is inaccurate but not greatly so given the evidence on the income elasticity of heating and given the scale of the fuel allowance. When expenditure and grant are integrated, then the relevant condition is either,

$$w'(x) = v'(-x + y) \quad (10.2)$$

or

$$w'(x) = u'(y - x) \quad (10.3)$$

The first case applies when heating expenditure exceeds the grant; the second case is relevant when the grant is greater than heating expenditure. Totally differentiating the first of these expressions with respect to y yields,

$$\frac{dx}{dy} = \frac{v''}{w'' + v''} \quad (10.4)$$

Differentiating the second expression produces.

$$\frac{dx}{dy} = \frac{u''}{w'' + u''} \quad (10.5)$$

The crucial issue here is the shape of the valuation functions. In the standard version of prospect theory, the second derivative of u is negative and we have assumed a similar property for w . As a result, when the grant is higher than heating expenditure a rise in its level increases heating. This is what we would anticipate. On the other hand, with the standard formulation of prospect theory, (as in Figure 10-1) v is convex. Thus the sign of 10.4 is ambiguous, but if in fact the first order conditions represent a maximum then the denominator of 10.4 is negative, implying that, with v'' positive, a rise in the heating allowance lowers expenditure on heating.

One way to 'fix' this problem is to drop the assumption that v is convex - a feature of the valuation function that makes some sense in a risky context, but perhaps less so in the absence of risk (see Chapter Two or Munro and Sugden, 2003, for more detail) On the other hand, mental accounting can also be integrated into other theories and this might represent the more straightforward option. For instance, one simple format⁴ is to write the utility function as $u = u(x_1 - \lambda_1 m_1/p_1, \dots, x_n - \lambda_n m_n/p_n)$ where $m_i \geq 0$, $i = 1, \dots, n$ represent labeled mental accounts, $p_i \geq 0$, $i = 1, \dots, n$ are prices and the $\lambda_i \geq 0$, $i = 1, \dots, n$ are coefficients representing the 'stickiness' or effectiveness of the label.⁵ In some cases $m_i = 0$ or $\lambda_i = 0$. A mental account of £1 for good i , is therefore akin to £ λ_i of committed expenditure. Figure 10-2 illustrates the effect of the introduction of a mental account of m_1 in the context of a two good model with prices set equal to 1. E_0 is the consumer's choice in the absence of the account. Setting the account at m_1 means the consumer treats $\lambda_1 m_1$ as a committed expenditure that makes no contribution to utility. Thus consumption of x_2 at the new equilibrium, E_1 is equal to that which would obtain if there was no mental account but the individual had $\lambda_1 m_1$ less income (i.e. at E'_1). Meanwhile consumption of x_1 is equal to $\lambda_1 m_1$ plus the value of x_1 at E'_1 . Letting $m = (m_1, \dots, m_n)$ we can therefore write the demand function as,

⁴A more general format is to suppose that a vector of mental accounts, m , represents a reference point, not in goods but in expenditure. A rise in any element of m , then favours bundles with higher levels of the corresponding consumption good. Imposing continuity of preference in m and homogeneity in prices and m together as further reasonable assumption leads to a generalisation of consumer theory (Munro, 2006b).

⁵This simple formulation therefore implies that demand is homogeneous of degree zero in mental accounts and prices taken together.




Figure 10-2 The Effect of a Mental Account on Choice

$$x_i(p, M, m) = \frac{\lambda_i m_i}{p_i} + x_i^*(p, M - \sum \frac{\lambda_i m_i}{p_i})$$

where x_i^* is a standard Marshallian demand function that depends only on prices, p , and an adjusted measure of income.⁶

The crucial feature which makes the policy more effective (in the sense of increasing the stickiness coefficient, λ_i) is the effective creation of the mental account. With the Winter Fuel payment, the name provides one stimulus to its creation. In addition, pensioners have to apply for the payment which may create some sense of obligation to spend the money received on heating. Finally it is heavily advertised and paid at a time of year (December) when fuel consumption is rising. As we saw with Prelec and Loewenstein, 1988, results, timing the payment to match the period of highest consumption during the year, may lower the psychic costs of expenditure.

Given the value of λ we can illustrate the effect of labelling a transfer through Figure 10-3. A rise in income of Δm pushes out the budget constraint and increases heating expenditure from x_0 to x_1 . If the same increase in income is labelled as heating expenditure and $\lambda_1 = 1$, then heating rises to x_2 . For $0 < \lambda_1 < 1$, heating expenditure will lie between x_1 and x_2 .

More formally, the responsiveness of expenditure to a marginal rise in the mental account is given by,

$$\frac{\lambda_i}{p_i} \left(1 - \frac{\partial x_i^*}{\partial M} \right) + \frac{\partial x_i^*}{\partial M}$$

So, as long as all other goods taken together are not inferior, an increase in income into a labelled account will therefore have a bigger increase in expenditure on the nominated good, compared to an unlabelled increase in income of the same size.

Winter Fuel Payments are not the only labelled accounts. In the UK the Budget for 2003 introduced the Government's proposals for its Child Trust Fund. Children born after September 2002 receive an endowment at birth,

⁶The formulation is reminiscent of the Stone-Geary or linear expenditure system or, more generally, the committed expenditures of the consumer demand model in Lewbel, 1985.




Figure 10-3 The Effect of Labelling on Heating.

means tested on parental income, but with a minimum value of 250 pounds and a maximum of 500 pounds. A top-up at age seven of £250 is also provided by the government. Until their 18th birthday, the child's grant is placed in a tax-free trust, but parents and others may add to it at a maximum rate of £1,200 per year. The impact and duration of the scheme remains to be seen and uptake has been patchy, but at least in standard theory, its effect is likely to be minimal: many families have at least some savings nominally put aside for matters such as college education, marriage expenses etc.

The impact of the Winter Fuel payment and Child Trust funds has not yet been clearly assessed. On the other hand, Child Benefit has been around a lot longer and also has large elements of the labelled mental account, particularly since in many countries where it is paid (e.g. the UK, Netherlands) there are no constraints on how the benefit should be used. Kooreman, 2000, using data for the Netherlands concludes that the marginal propensity to spend on child clothing out of child benefit is 10 times that out of ordinary income. However, Edmonds, 2002 (for Slovenia) and Blow et al, 2005, (for the UK) find no distinctive impact of child benefit on expenditure. Indeed the latter find only adult alcohol consumption to be distinctively sensitive to child benefit changes.⁷ What makes the U.K. so different from the Netherlands is not clear, but it might lie in the way the benefit system labels and maintains the child benefit 'mental account'. In the UK payments are monthly while in the Netherlands, Child Benefit is paid quarterly making it larger and possibly more salient in the minds of parents.

In these examples accounts are created by naming and advertising as such, but other fiscal policies may create a mental account as a byproduct, with the nature of it sensitive to the framing of the policy. For instance, in the Hickian model a subsidy for saving can also be implemented as a tax on current consumption accompanied by a lump-sum rebate. These two policies could have quite different effects on the creation of mental accounts and therefore on behaviour.

⁷Lyssiotou, 2005, uses a demand system rather than individual equations and, using the same data as Blow et al, 2005 does find a significant labelling effect for the UK - suggesting that differences in econometric methods might lie behind some of the disagreement.

Promoting heating use by UK pensioners might be viewed as a sensible merit-wants style policy, but in theory governments can use mental accounts for other purposes. For instance in the theory of Leviathan, the government sets its tax policy so as maximize the feasible size of government. In doing so it faces the constraint of obtaining the consent of the electorate. Having many small taxes may help obscure the size of the government's income, but in addition, hypothecation may reduce the perceived deadweight loss from taxation, if citizens integrate the tax payment with the benefits from the expenditure side.

Overall, therefore, creating mental accounts by labels and by the appropriate timing of payments may be a means to raise the income effects associated with some forms of government expenditure. However, as the empirical example of Child Benefit in the UK suggests, there is a need for some scepticism about the universal effectiveness of this kind of policy.

10.3 Intertemporal issues: pensions and savings policy.

There are two different aspects to bounded rationality in the context of decisions about savings and pensions. The first concerns the intra-temporal allocation of invested funds; the second and the most prominent concerns the inter-temporal distribution of consumption. Rapid changes in our understanding of inter and intra-temporal choices have been fuelled by the emergence of behavioural finance as an alternative to the rational choice model of financial behaviour (Schleifer, 1999). In part the growth of the behavioural approach has been prompted by some of the anomalies discussed in Chapter Two and their persistence as anomalies despite many attempts to bring them within the scope of the rational choice model. Schefrin, 1999, lists examples of regularities which seem to defy explanation by the rational choice model, including the equity-premium puzzle (whereby, over long periods of time the return on stocks exceeds the returns on bonds by an amount which is not consistent with standard values for individual risk aversion), and the tendency to hold on to losing stocks too long and sell rising stocks too quickly (Odean, 1998).

The second aspect of bounded rationality concerns the inter-temporal allocation of consumption, where there is some evidence that consumers in many countries (particularly the Anglo-Saxon economies) undersave, relative to some desirable standard (see Bernheim, 1993 for the USA and Börsch-Supan and Brugiavini, 2001, for summary of the more heterogeneous European position).. As a result many households face sharp drops in their standard of living upon retirement. For instance, Banks et al, 1998, find a significant fall in the consumption of non-durables by UK households around retirement age. After controlling for some possible explanations, including mortality risk labour market participation and job-related expenses, two thirds of the drop is explained. For the remainder they conclude that they are only the sudden arrival of a large (and negative) shock to wealth can reconcile the observed behaviour with the life-cycle model of consumption. Similar patterns have been found for Canada by Robb, et al 1989, and for the USA by Bernheim et al, 2001 where consumption declines of up to 30% are found amongst the

households with the lowest wealth.

Now, most individuals who save have multiple opportunities to invest and to observe the consequences of their previous choices. So if there is any instance where learning might supplant the need for government control, then perhaps intertemporal choice is it. Nevertheless a consensus that many individuals undersave suggests that the forces which create the apparent undersaving are quite powerful and distinct from some of the anomalies discussed earlier in the book. In the next sub-section I discuss some of the evidence on intertemporal preferences, together with theories which attempt to explain the evidence or to build models of investor behaviour on the back of it. Thereafter we return to the discussion of policy instruments.

10.3.1 Intertemporal choice and self-control.

The two prominent problems facing the rational choice model in the context of intertemporal choice concern, firstly, giving the wrong weight to future selves when planning consumption (i.e. a failure of strong rationality) and secondly, the issue of dynamic inconsistency (a failure of weak rationality), meaning that the optimal plans made in one period may not be optimal when viewed by the selves of other periods. To a significant degree, behavioural economics has joined these two concerns by picturing the typical consumer as someone who overweights the present because of intertemporally inconsistent preferences.

Suppose preferences are additively separable between time periods and can be represented by a utility function of the form,

$$U_t = \sum_{\tau=0}^{\tau=T} \delta_{t+\tau} u(x_{t+\tau}) \quad (10.6)$$

where $x_{t+\tau}$ is the consumption vector in period $t + \tau$ and $\delta_{t+\tau}$ is the discount factor attached to consumption in that period. Meanwhile suppose that individual welfare can be written as,

$$W_t = \sum_{\tau=0}^{\tau=T} \delta_{t+\tau}^* u(x_{t+\tau}) \quad (10.7)$$

Then the strong rationality issue can be defined as $\delta_{t+\tau} \neq \delta_{t+\tau}^*$ for some $\tau \geq 0$. Clearly, if it were for not for the time interpretation subscripts, then the wrong weights attached to the utility for individual periods would be seen as simply a specific case of merit wants. The time subscript is important because of the intuition it provides about the direction of bias: it is usually supposed that $\delta_t > \delta_t^*$, in other words the current period is overweighted in deliberations, compared to the interests of future selves. This view has a long history in economics as Peart, 2000, points out. Stanley Jevons, for instance stated in the Theory of Political Economy that

‘to secure a maximum benefit in life, all future pleasures or pains, should act upon us with the same force as if they were present, allowances being made for their uncertainty’ (1871, page 72).

He contrasted this ideal with the actual state of affairs, ‘a future feeling is always less influential than a present one,’ (1871, page 72). Alfred Marshall and Irving Fisher expressed similar views, with the latter stressing that poverty made the distortions to preference that much greater. ‘The smaller the income a man has, the higher is his preference for present over future income’, (Fisher, 1910, page 381). But the most well-known precedent for the view that current consumption is overweighted compared to future needs is associated with the name of Pigou, who argued that a failure to give equal weight to future consumption means only ‘that our telescopic faculty is defective and that we therefore see future pleasures, as it were, on a diminished scale’ (Pigou, 1920, page 25).

The second feature of bounded rationality in the context of intertemporal decision-making concerns the failure of future selves to necessarily follow the plan of consumption and saving laid out by the current self. Let X_t be the choice set for the individual at time t , a typical element of which is $x = (x_t, \dots, x_T)$

Define a plan as,

$$x^*(t) = \arg \max_{x_t, x_{t+1}, \dots \in X_t} \sum_{\tau=0}^{\tau=T} \delta_{t+\tau} u(x_{t+\tau}) \quad (10.8)$$

Write $x_{t+1}^*(t)$ to mean $x^*(t)$ excluding x_t , then dynamic consistency is the property that, for all t , $x_{t+1}^*(t) = x^*(t+1)$. In other words, future consumption planned at time t is also optimal at later times. As Strotz, 1955, first pointed out, a sufficient condition for dynamic consistency is a constant rate of discounting, that is $\delta_{t+\tau} = \delta^\tau \delta_t$ for some $\delta \geq 0$ and all $t, \tau > 0$.⁸

The evidence that preferences do not have this simple feature is widespread, but not uncontested. A recent survey can be found in Frederick et al, 2002. In experimental work, discount rates are typically elicited by asking subjects to choose between x now and $x+y$ t periods into the future. Alternatively subjects are asked for the value of y which makes them indifferent between the two options. The per-period discount rate, δ is then calculated as the value which solves the equation, $x = (x+y)\delta^t$. In an early, but fairly typical study, Thaler, 1981, finds a value for δ of 0.22 (equivalent to a discount rate of 345%) when t represents one month, $\delta = 0.45$ for a one year interval and $\delta = 0.84$ when t is equal to 10 years. In other words, δ rises as t increases, rather than being constant as dynamic consistency would require. This pattern is usually referred to as ‘hyperbolic discounting’ on the grounds that a hyperbolic function often fits the data on δ_t fairly well, but in fact it seems more likely that the current period discount rate, δ_1 is low relative to discount rates in subsequent periods which thereafter remain fairly constant.⁹ As a result, the following form for the evolution of δ has been proposed:

$$\begin{cases} 1 & t = 0 \\ \beta\delta^t & t > 0 \end{cases}$$

⁸It is also necessary for most of the budget sets commonly encountered in models of intertemporal choice.

⁹Read, 2002, argues that discounting is sub-additive rather than hyperbolic. In other words, when time periods are sub-divided, the sum of the resulting discount rates exceeds that for the undivided period. This is similar to the part-whole bias found in the valuation of goods (see Bateman et al, 1997a for instance).

where $\beta < 1$ (see Laibson, 1997). More recently, Harrison et al, 2003, have used field and laboratory experiments to challenge much of the accepted wisdom on hyperbolic discounting. They note that typically subjects are given the choice between receiving payment now, in the laboratory or at a later date. Delaying payment introduces some uncertainty about whether payment will be received at all. Delaying all payment controls for this effect and, on the basis of their Danish evidence reduces the rise in δ over time, though without altogether eliminating the evidence for hyperbolic discounting in humans.

In addition to the evidence on hyperbolic discounting there is also some interesting experimental work (though much of it involving hypothetical choices only) which questions a number of the other assumptions of the basic intertemporal model. Prelec and Loewenstein, 1992, for instance, observe that typically individuals have some preference for intertemporal variety, thus violating additivity. Faced with eating 100 meals of chicken tikka in 365 days, many individuals would prefer to spread them out evenly over the year, rather than eating them sequentially. They also show that individuals prefer an increasing sequence of payments (e.g. £1,000, £2,000, . . . £5,000) to the same sequence in reverse, even though the latter dominates the former given positive market interest rates.

An optimal consumption plan that is not dynamically consistent raises the question of what the actual plan will be. Two factors are important in this context:

- The agent's ability to commit themselves to a future plan of consumption;
- Their degree of awareness of the dynamic inconsistency.

The individual who is not clear-sighted may stumble from plan to plan. In each period she or he plans consumption for all future periods and, in each subsequent period plans only recently made are torn up and a fresh plan of lifetime consumption is laid out. This combination of a long planning horizon and myopia about the actions of their future selves lacks psychological credibility. At the other extreme, and possibly equally incredible, is the clear sighted agent who anticipates the decisions of future selves and hence is constrained to choose from only those plans which are dynamically consistent.

One important strand of the extensive literature on dynamic inconsistency has tried to say something about welfare. To a degree this means adopting the perspective of one of the life-time selves involved in the planning of consumption. For instance in the over-analysed instance of Odysseus and the Sirens, the optimality of his plan to lash himself to the mast is based on giving greater weight to the views of the pre- and post-Siren Odysseus compared to his preferences as he passes the rock. Most people will probably not object to this judgement in the case of Odysseus, but in other cases of dynamic inconsistency it is much harder to settle on a clear-cut view of the correct metric of welfare. As such, an optimal public policy may not be defined for dynamic inconsistency in general.

For some models of intertemporal choice, the welfare issues are clearer. An instance is provided by the theory of procrastination advanced by O'Donoghue and Rabin, 2001. They consider an infinitely-lived individual who must decide whether and when to undertake a task which will cost c in the current period, but then yield a stream of benefits v in all future periods. In this

context procrastination is defined as never doing a task which has a positive net present value. The individual's future consumption is discounted at a rate of δ per period. It is supposed that the individual also underweights future consumption – placing a weight of $\beta \leq 1$ on all benefits or costs which accrue beyond the current period. O'Donoghue and Rabin distinguish between the preference for immediate gratification (as indicated by β) which leads to time-inconsistency and the subject's naivete about their future value of β . A subject who correctly anticipates their future value of β is described as sophisticated, while someone who believes that their future value of $\beta = 1$ is completely naïve, unless β is actually equal to one. I have already noted that complete naivety is not credible in a context where agents are making lifetime plans, but O'Donoghue and Rabin argue that models of dynamic inconsistency are not robust to the introduction of some naivety (i.e. a forecasted value of β somewhere between its true value and 1). O'Donoghue and Rabin advance the notion of a 'perception perfect strategy', a concept akin to that of subgame perfection, in which a) subjects choose optimally in each period, given their beliefs about the actions of their future self and b) their beliefs about their future beliefs are correct. Note though that what an individual actually does in any given period may contradict what they believed they would do in previous periods.

An example illustrates their model and its main results. Suppose, in task A, $c = 5, v = 0.8$ and $\delta = 0.9$. The present value of doing the task in the first period is $-c + \delta v / (1 - \delta)$, which is 2.2 in this case. However, the individual actually appraises the value of the task as $-c + \beta \delta v / (1 - \delta)$, which for $\beta = 0.8$, makes the value of the task equal to 0.76. So the task is still worthwhile, in the sense that it has a perceived positive net present value. However, its net present value as seen by the agent, may be still higher if implementation is delayed by one or more periods, because this has a more marked impact on the value of the cost of the project, compared to the stream of future benefits. Implementation in the t th period yields a payoff of $\beta \delta^t [-c + \delta v / (1 - \delta)]$. Figure 10-4 shows the path of net present values for implementation of the task in periods one to ten. As can be seen, with sophisticated beliefs, the individual prefers any delay of up to eight periods, compared to implementation of the task immediately. However, doing the task in period one is preferred to leaving it until period nine or later. Suppose, in period 1, the individual believes that future selves will not implement the plan in periods 2 to 8, then he or she will prefer to do the task now rather than delay. Hence the longest possible delay is eight periods and all equilibria involve the agent executing the task in one of the first eight periods.

Now consider an individual who is naïve about the beliefs of future selves, thinking that future selves will weight future consumption by $\beta^* = 0.9$ rather than $\beta = 0.8$. The upper time path in Figure 10-4 shows the path of net present values for a sophisticated person with $\beta = 0.9$. This is how the partially naïve person believes futures selves will behave. In other words, futures selves will not tolerate a delay of more than one period. The partially naïve person then compares doing the task immediately to leaving it, in which case, they believe, the maximum future delay will be one more period. Since the discounted value of executing the task in period three is higher than its value in period one, they choose to postpone doing the task. Their perspective and analysis is repeated in the next period and so the task is never completed.

The critical factor here is that the maximum tolerable delay for the naïve

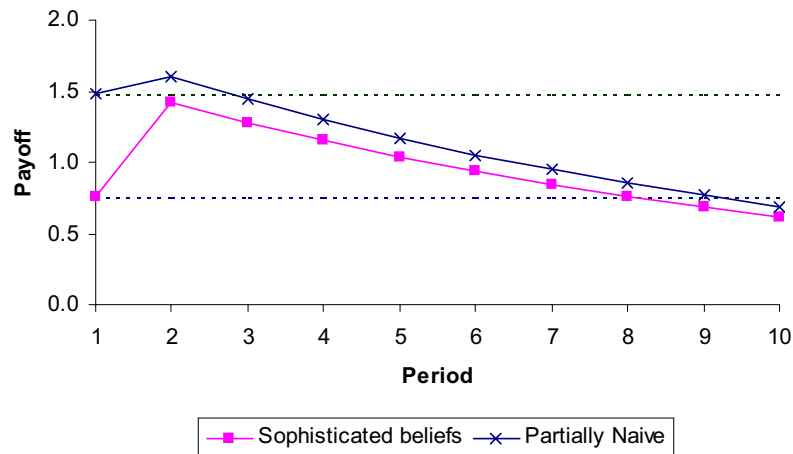


Figure 10-4 The Value of Procrastination.

individual is strictly shorter than its equivalent for the sophisticated individual. Whenever this occurs the task will remain undone.

O'Donoghue and Rabin go on to show that procrastination may occur with a larger choice set when it does not occur on a smaller sub-set of choices. In other words, increasing the choices open to an individual can induce procrastination. For instance if the individual could only choose between doing nothing and implementing task B where $c = 1$ and $v = 0.3$, then for the same values of β and δ used above, the individual would implement the task immediately, whether they were naïve or not. Adding the option of doing task A used above means that if the individual were to execute a task then it would be A, since it has a higher net present value. However, as we have just seen, a sufficiently naïve individual would never do task A.

As the authors point out, the model has its limitations: the naïve individuals never learn from their faulty forecasts of future behaviour and neither they nor their sophisticated brethren employ any of the psychological and incentive devices that many individuals employ to avoid extensive delay of the kind the model produces. The model though does illuminate a number of issues. Clearly, complete procrastination of the kind found in the model is welfare-reducing for the individuals concerned. However, some delay often raises perceived utility, hence implementation in the first period may be against the perceived interest of some selves. As with Odysseus therefore commitment devices can have value.

The other feature of the model which is of interest is methodological: the agents mix quite sophisticated reasoning with, in some cases, extreme foolishness. The agents are assumed to be able to evaluate an infinite stream of payoffs and compare implementation of the task in different periods. They are able to posit the existence of future selves with different preferences. They select only plans which are consistent with their beliefs about the actions of these future selves. On the other hand, naïve agents underestimate their own potential for future procrastination and they fail to learn from the failure of their earlier plans.

Many people seem to be aware of the procrastinating feature of their own decision-making and make commitments to try to overcome procrastination. Mayer and Russell, 2005, offer the example of teachers in some districts of Cal-

ifornia many of whom choose to receive their income in 12 annual instalments when the possibility of getting it in 10 payments is also available. For an older example, which uses a different language for intra-personal conflict, we turn to Charles Madge, 1941, who in his investigative tour of wartime saving in English cities noted the large numbers of Christmas clubs and saving schemes:

In psychological jargon, it is the function of the super-ego to make people save – not many are natural misers. Therefore, because of the known weakness of the super-ego in face of temptation, Lancashire women mill-workers are prepared to pay somebody to come round once a week and collect their sixpences or shillings. Page 429-430.

Given these observations on intertemporal anomalies we turn to the issue of savings policy.

10.4 Regulating the framing of savings and pensions choices.

In the UK in 1988, the then Conservative government liberalised the pensions market (Connolly and Munro, 1999). Previous to the reforms, individuals had been heavily constrained in their choice of options. All employees with earnings above a minimal level were forced to make contributions towards the first tier of the pension scheme – the Basic pension, the level of which was not income related. For the second tier of the pension system – the part related to income - employees could choose between the state scheme (SERPS) or could join an institutional scheme organised around their workplace. Typically, members of such schemes paid a proportion of their income each month into the scheme, as did their employers. As an inducement National Insurance contributions were reduced for members of such schemes.

The reforms left the structure of the first tier largely intact, but it added an extra element of choice for the second tier. Now, employees could opt for a pension scheme provided by a private supplier, provided that such a scheme satisfied certain minimal requirements. The advantage of the private schemes lay in their portability. Typically, in pension schemes organised through an employer, an individual who left before retirement faced reduced benefits compared to individuals who maintained a single employer throughout their working life. For individuals who switched jobs regularly or who had a pattern of interruptions to employment, private schemes could be more advantageous because they were attached to the individual rather than the company. On the other hand they did have some clear disadvantages for many potential purchasers (Blake 1995). First, administrative costs were typically higher; secondly employers were typically not as generous in their contributions to private saving as they were to the collective pension schemes. Finally, almost all private schemes were defined contribution (DC) schemes, meaning that the value of the pension was determined by, for instance, the values of shares and stocks at the time of an individual's retirement. Most employer-organised schemes, on the other hand, were defined benefit (DB) schemes, meaning that

the generosity of the pension was some fixed fraction of an averaged measure of an individual's lifetime earnings. In short therefore, risks for workers were higher with DC schemes, compared to their more traditional options.

Despite these apparent deficiencies, private schemes were immediately popular. By the early 1990s several million individuals had taken out private pensions. In many instances, individuals switched out of employer-organised DC schemes into private pensions. What was startling about the up-take however, was not its level but the individuals who switched – hundreds of thousands of state-employees such as nurses and teachers opted out of their existing scheme and into the new pensions (Disney and Whitehouse, 1992). Nurses and teachers in the UK typically work for the state throughout their lifetime, their jobs are relatively secure and employer contributions to the pensions scheme are relatively generous. Hence, they were definitely not the people for whom the reforms were originally devised.

It is by no means certain that all these individual choices were irrational. For some people private pensions released them from a rigid pattern of savings which dictated that individuals made heavy contributions to pensions at times in their lives when, perhaps, they would have preferred to spend money in other ways. That said, it was clear that many people rapidly regretted their decisions and attributed them to heavy pressure from salespeople combined with misleading information about the consequences of their decisions. Above all there was the temptation to put off saving for the future and to enjoy a temporary rise in consumption. Much energy in the UK pensions system since the early 1990s has been devoted to sorting out the aftermath of the reforms, reversing hasty choices to opt-out and compensating individuals who claimed they were misled.

The example of pension policy reform suggests some of the forces which tend to limit saving for pensions:

- Having to opt-in to higher savings levels
- Lack of clear and reliable information on the impact of savings rates on future consumption.
- Having a 'savings' account or policy rather than separate policies for different forms of saving.
- Having the possibility of raising saving in the future, but without commitment.
- Immediate consumption benefits from reductions in saving.

Not surprisingly therefore, Thaler, 1994, argues that the ideal savings programme should have four features (see also Weiss, 1991):

1. it provides an immediate reward to saving
2. it is simple to use and understand
3. the money is perceived as 'off limits' to current spending
4. experts consider the programme a good deal. (Thaler, 1994, p. 189)

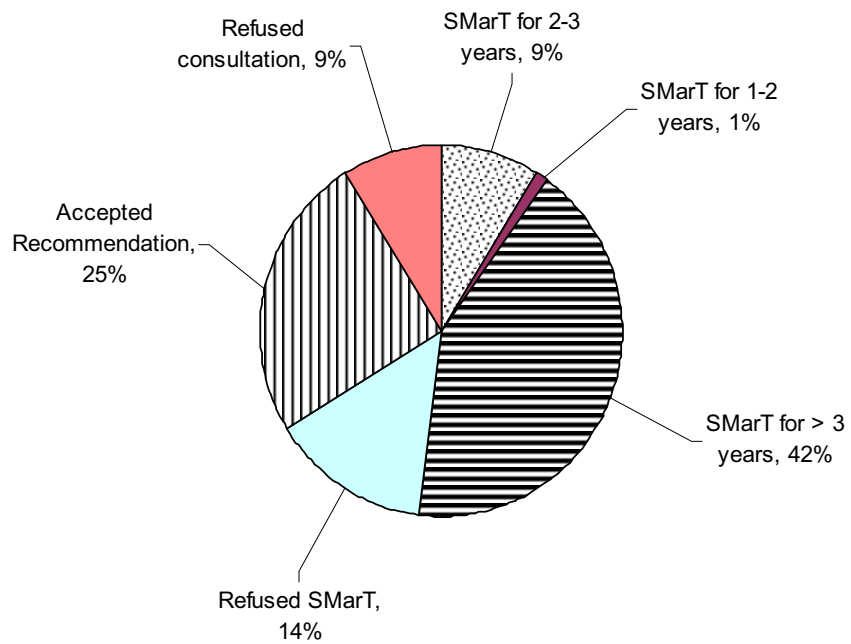


Figure 10-5 Choices in Thaler and Benartzi, 2004.

Thaler and Benartzi, 2004, offer an example of the theory at work, involving a US company which wished to raise the savings rates of 315 of its lower paid employees. All were offered the chance to talk over their retirement planning with a consultant. Of those who accepted the offer, around one quarter then took the initial advice of the consultant. For the remainder, Thaler and Benartzi had devised a savings plan (SMarT) in which members would make no immediate rise in their contributions but would agree to commit 3 percentage points of their annual wage rises to their retirement savings plans. Subjects could, if they wished, change their minds at a later date and withdraw from their new commitment.

Figure 10-5 shows the results of the exercise. Thirty four percent of the group either refused the opportunity for a financial consultation or accepted the counsellor's suggested change of savings behaviour. About one in seven then refused to enter SMarT, but of the remaining 52 percent of individual, nearly all then stuck with the programme for three years of pay awards. For those who joined the scheme, average savings rates were 11.6 percent after four years, compared to the 3.1 percent prior to its introduction. In fact after three years, their savings rates exceeded those for the individuals who had taken the consultant's original advice. Meanwhile members of the workforce who had remained outside the scheme had savings rates which hovered between six and seven percent.

The theory behind this scheme is closely related to the mental accounting literature discussed in Chapter Two coupled with the problems of intertemporal optimisation discussed in the previous section. The source of the problem is viewed as procrastination on the part of the workforce. An immediate rise in savings rates, even when offered, is not taken because the pain of immedi-

ate loss of consumption is hard to bear given only the promise of the pleasure higher consumption in the distant future. But if as we saw, the pain of consumption foregone is lower than that from a reduction of actual consumption, individuals will be more willing to accept a scheme such as SMarT over an immediate rise in contributions.

Although the positive side of SMarT-like policies is relatively clear, there still remains an important normative issue. The Thaler scheme passes the test of ex ante and ex post consent, because quite clearly participants can opt not to join or can opt out of the savings programme at a later date. As a result they also pass the related ‘libertarian paternalism’ test discussed in Chapter Five - fully rational individuals would not find their path to optimisation hindered by the scheme, but boundedly rational agents appear to benefit.

Following the Thaler guidelines, Gorrige, 1998, draws out the following general sorts of policy for governments which wish to see higher rates of saving for retirement.

- Legislate to make joining a (private) pension scheme a default option for everyone.
- Increase withholding tax rates
- Make the default option for tax refunds their payment into a pension scheme or other contractual savings account
- Provide information on saving for retirement
- Subsidies for saving

The last two elements are fairly standard devices. Each of the first three is designed to create inertia around the decision to save rather than consume. If taxes are withheld from income, then this limits consumption, but if the eventual refunds are then sent into savings accounts then this further reinforces the default option of saving rather than consumption.

10.4.1 Planning ahead.

One of the features of savings activity is its variability even after controlling for standard economic and demographic variables. In other words not everyone seems to suffer from problems of self-control. Across all ages and all income classes, some people save while others do not. For instance, for the U.K. Banks et al, 2002, show that around 10% of households in the top half of the income distribution have no financial wealth. Even in the richest 10% of the population, one in twelve households has £500 or less in financial assets. However, this diversity should not blind us to the fact that it is principally low income households that lack savings. Over half of the poorest 10% of households have no savings, while Banks et al, 2002, point out that for the poorest 20% of the population the median savings is just £50.

Psychologists (e.g. Warneryd, 1980) have examined the traits and attitudes that make some people save while others in similar economic circumstances

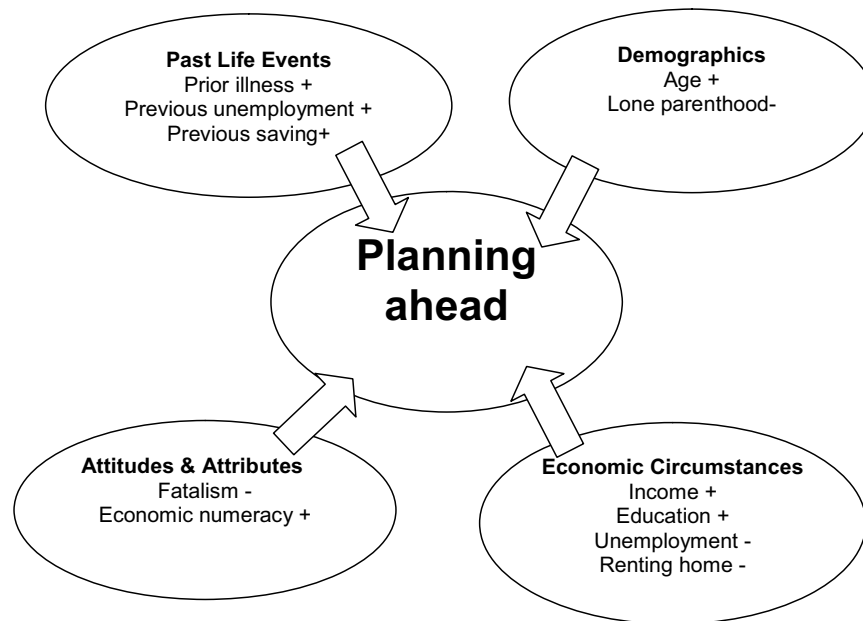


Figure 10-6 Factors Affecting Planning Ahead.

do not. Not surprisingly, conscientiousness is an important factor associated¹⁰ with saving (Webley et al, 2001) along with the personal discount rate, which tends to decrease as income and age rises. Meanwhile when Lunt and Livingston, 1991, examined the lives of a group of Oxford residents, they found that against the positive influence of education, income and age on saving, non-savers have greater preference for flexibility, tended to be more fatalistic about the future, were less likely to monitor their expenditure or to plan ahead and were more secretive about finances with friends and relations.

In the language of Ajzen, 1991, much saving is planned behaviour which is most likely to occur when personal attitudes towards saving, beliefs about its efficacy and feelings of social pressure are aligned. Effective saving behaviour requires a person to plan ahead and to believe that setting aside money will make a difference. Not surprisingly therefore, 'planning ahead' is a key predictor of whether people save. In turn, planning ahead is correlated with a large number of factors which are summarised in Figure 10-6. (The signs in this figure refer to whether the factor has a positive or negative correlation with planning.)¹¹

An influence on planning ahead and a theme running through much of the work on savings psychology is the notion of 'locus of control', a concept from social psychology (e.g. Lefcourt, 1982) which refers to the extent to which

¹⁰The issue of causation rather than correlation has not been a central feature of this literature.

¹¹Some sources: prior illness, previous employment (Atkinson, 2006), previous saving, age (Webley, et al 2001), lone parenthood (Atkinson et al, 2006), fatalism (Lunt and Livingston, 1991, Eckel et al, 2001, Wu, 2005), economic numeracy (Banks and Oldfield, 2005), income (Lunt and Livingston, 1991, Atkinson et al, 2006), education (Lunt and Livingston, 1991), unemployment (Atkinson et al, 2006), renting home (Atkinson et al 2006).

individuals feel in control of what happens in their life. Individuals with an internal locus are people who feel in control; individuals with an external locus are more fatalistic, believing that events outside of them determine what goes on in their lives. Its relevance here is that non-saving behaviour and not planning ahead is associated with an external locus. Direct evidence of this link is provided by an important savings experiment financed by Human Resources Development Canada (Eckel et al, 2001). A wide range of subjects participated of whom, 72 per cent had a family income of less than 120% of an official Canadian poverty line. In one treatment, participants were offered a real choice between 100 Canadian dollars in cash one week after the experiment or an investment of 400 dollars in a fund for education and training available for spending several months later. Subjects who had revealed high discount rates in another part of the experiment were more likely to opt for the immediate cash, as were subjects with an external locus of control.

In a similar vein Wu, 2005, uses the 4,400 respondents to the 2001 US Survey of Consumer Finances to explore why people save. Out of the sample, 18% do not save at all. Those realising that they need to save more are more likely to be pessimistic about the future of the economy and more likely to agree with the view that they are unlucky with their finances. However this group is less likely to have done something about it and when they do save, they are less likely to shop around for financial products.

The ability to plan ahead is also affected by maths skills. Using the 2002 wave of the English Longitudinal Study of Ageing, Banks and Oldfield, 2006, conclude that amongst the over 50s, a large proportion of the population have poor financial numeracy skills. The ability to solve relatively simple exercises in financial arithmetic is correlated with knowledge and understanding of financial products and also with levels of saving and retirement provision, even after other factors are controlled for¹². The results match those of a wider study concerning the financial capability of the UK population, carried out on behalf of the Financial Services Authority, Atkinson et al, 2006, find a number of factors positively linked with planning ahead, including education, age, income and a number of factors negatively linked including renting, being unemployed, being a lone parent, and not having a current account. Those planning ahead were also more likely to have been active choosers of financial products, to keep informed about money and to have a high score in relatively simple exercise in financial arithmetic.

Inferring causation from some of these associations is difficult. With the exception of the Eckel et al, 2001, study, the results indicate correlation rather than causation. Nevertheless, there have been some attempts to develop savings policy for low income earners on the assumption that financial training and heavily subsidised, small-scale saving can reduce fatalistic attitudes and encourage greater planning ahead. The most notable attempts to innovate can be seen in the case of anti-poverty policies designed to raise wealth rather than income.

The notion of asset-based welfare and particular, Individual Development Accounts (IDAs) has provided a radical challenge to the traditional format of the welfare state, given that historically, welfare minded governments have focused on providing support for incomes rather than savings. As conceived by Michael Sherraden, 1991, IDAs were part of a proposal to reverse the

¹²Bernheim and Garret, 2003, report on the positive effects for saving of workplace-based financial education.

emphasis and make asset based policies a central component of the welfare state. To that end IDAs consist of three main components: a matched subsidy to saving for low income groups; education in financial management skills for the recipients and restrictions on the use of the subsidized savings for uses such as homeownership or education. Though still experimental, IDAs have become a popular bipartisan policy in the US and to date over 40 states have approved them in some form or another, usually delivered with the active participation of the not-for-profit sector. With the introduction of the Child Trust fund scheme (mentioned above) and the development of the Savings Gateway project, asset-based welfare has spread to the UK.

Most IDAs are non-experimental: community groups and municipalities offer themselves as possible sponsors and then actively recruit potential participants. As such scientific assessment of their impact has been limited. The most compelling evaluation of IDAs is that conducted by Miller et al, 2004, reporting on an IDA in Minnesota, typical except for the fact that eligible applicants were randomly accepted or rejected for the programme. Fifty-three per cent of people who signed up for the programme closed their account without ever qualifying for matching funds. Out of those who remained, average savings were \$1,480 per participant, compared to \$2,041 total financial assets (including retirement savings) at the start of the scheme. The largest long term impact of the scheme was on home ownership where after 48 months (1 year after the programme ended) those randomly accepted into the IDA had ownership rates 6 percentage points higher than those rejected. The wider impact of a group of 14 IDAs, linked under the heading of the American Dream Demonstration (ADD) is assessed in Moore et al, 2001. They report participants feeling higher confidence in the future (93% of their sample) and feeling more in control over their lives (85%).

The Savings Gateway programme has been the UK's pilot version of asset-based welfare. In the initial pilot (there is now a second) subjects were recruited non-randomly into five different schemes based in low-income urban districts. As with the U.S. schemes, savings were heavily subsidized and, with the exception of one variant, accompanied by some financial education. In their evaluation of the first group of pilots for the UK's Savings Gateway programme Kempson et al, 2005, find that average participants save £361 out of the possible £375 over the lifetime of the projects. Perhaps more importantly, compared to a reference group of individuals from areas neighbouring the Savings Gateway pilots there is some evidence of a positive change in attitudes towards saving – 39% of the survey sample agreed or strongly agreed that they felt more in control of their life as a result of the scheme, while 60% felt more financially secure and a third stated that they were more likely to save for retirement. One of the schemes, in Hull, had no financial education and targeted potential recruits using benefit data, without the aid of community groups as intermediaries. As such, the Hull scheme represents a lower cost approach to asset-based welfare, compared to the US style of IDA. The Hull group saved as much as the groups where community groups were more heavily involved. However, the large number of individuals who saved at or close to the maximum in all the pilot areas limit means that any underlying differences in the effectiveness of the schemes may have been masked by the low level of the cap on savings.

10.5 The framing of tax policies.

The idea that behaviour will depend on the framing of tax policies follows naturally from the arguments of this book, but it is important to distinguish two separate routes by which framing effects can be relevant for tax policy.

1. Since market behaviour is susceptible to a variety of framing effects, it would be no surprise if individuals were not also influenced by the framing of taxes and benefits.
2. Secondly, perceptions of tax fairness might well be susceptible to the manner in which the tax and benefit system is portrayed. As a result, a tax system might be politically acceptable in one frame, but unacceptable to citizens when presented in an alternative manner.

It is possible that these two routes are linked, but it is not a necessity. An individual with other regarding preferences could be immune from framing effects in their own choices, but nevertheless perceive that other citizens might not be so clear-headed in their assessment of government policies. As a result she or he would not be affected in their market choices by the framing of taxes, but could still have clear preferences over potential frames.

Schelling's, 1981 example illustrates how framing may affect perceptions of the most acceptable tax system. Consider this slightly modified version of the pairs (A versus B and C versus D) of choices faced by his students. In both pairs the choice is whether to target a higher subsidy for having children to higher income families or to families on low incomes. However in the two choices, the framing of the subsidy differs.

Option	For		High Income		Low Income
A	having children	chil-	£50,000 income families receive a £2,000 bonus	income	£15,000 income families receive £1,000
B	having children	chil-	£50,000 income families receive a £1,000 bonus	income	£15,000 income families receive £2,000
C	not having children	having	£52,000 income families receive a £2,000 penalty	income	£17,000 income families receive a penalty of £1,000
D	not having children	having	£52,000 income families receive a £1,000 penalty	income	£17,000 income families receive a penalty of £2,000

Since every system with bonuses for some individuals is equivalent to one in which there are penalties for the complementary individuals, the only consistent answers to this pair of questions are the pairs {A,C} and {B,D}, but in fact the majority of Schelling's class opted for {B, C}. Such political preference reversals may occur either because the individual subjects were themselves confused by the framing devices or that they considered others to be fooled. There is no evidence on which interpretation is the more correct.

Experimental tests of Schelling's example have produced mixed results to date. Traub, 1999, did not find support for Schelling's hypothesis in his data, but in dedicated experiments on the issue with 49 subjects, McCaffrey and Baron, 2004, discovered that subjects typically rated the fairness of bonuses

to low income couples more highly than bonuses to high income couples but switched ratings when the choices were presented in the penalties format.

Traub, 1995, used employees of 5 north German companies to examine wider issues of the dependence of perceptions of fairness on endowments. In one treatment, subjects were given information about a proposed tax schedule for a single, childless person then asked to propose a fair level of tax for a couple with 2 children and the same gross income. Some weeks later subjects were presented with information on a proposed tax schedule for the couple-with-children family in which the tax was at the level they themselves had proposed in the earlier questionnaire. They were then asked to propose a fair tax on a single, childless person with the same gross income. In other treatments the order was reversed. Traub found that in general tax fairness ‘indifference curves’ were non-reversible: the tax reduction for the couples with children in the first treatment was larger than the surcharge for a single person without children in the second treatment. In other words, people were generous in subsidizing the ‘desirable’ or ‘deserving’ but reluctantly to penalise heavily the relatively ‘undeserving’.

In surveying the general issue of tax framing on behaviour, Krishna and Slemrod, 2003 put forward a number of conjectures about the potential form of framing effects:.

1. Endowment effects – as with other choices, citizens may be averse to changes in the status quo.
2. Aggregation and splitting of gains and losses. According to Tversky and Kahneman, 1979, (see also earlier in this chapter), the asymmetric shapes of value functions for gains and losses suggest that governments could raise the acceptability of expenditure by splitting it as finely as possible, while reducing the unacceptability of taxes by aggregating losses. Conversely, Krishna and Slemrod, 2003, point out there may also be threshold effects – small gains and losses might not be noticed by individuals, raising the acceptability of many small taxes and reducing the attractiveness of splitting benefits into many fragments.
3. Timing effects. The timing of taxation and benefits might influence its acceptability. For instance, one reason for paying out the Working Tax Credit¹³ on a frequent basis is to reinforce the perceived gains from working. An annual lump-sum might be hard to manage for low income families, but might also lead to dislocation between the tax refund and the work that generated it. also point to a marketing literature which concludes that many consumers feel more favourably towards purchases when prices are expressed in ‘pennies per day’ rather in a lump-sum format. Conversely benefits become more attractive when expressed in overall lump-sums rather compared to when they are broken down into daily amounts. Krishna and Slemrod, 2003, suggest that politicians may seek to exploit such effects when describing the benefits and costs of their manifestors. Note though that the effect, if true, contradicts the aggregation argument given above. In fact it is only reconcilable in the presence of threshold effects.

¹³Working Tax Credit, formerly Working Family Tax Credit is the UK’s system for subsidizing work for low income families. Payments are made weekly or monthly (at the choice of participants) directly into bank accounts.

4. Obfuscation. McCaffrey and Baron, 2005, argue that reducing the transparency of taxes can raise their public acceptability. Because individuals value more highly what is brought to their notice, losses and gains may be hidden by simply avoiding placing them in the spotlight of consumer attention. Indirect effects of taxes, such their consequences for prices may not therefore be salient. In some cases, this may not be simply an issue of consumers being misled, but an active preference by citizens for hidden taxes. McCaffrey and Baron, 2005, cite the example of social security payments. If these are wholly framed as taxes from income then, they argue, workers will feel the pain of making the payments. However, if the deductions are defined as employer payments, then employees will not code the contributions as losses and therefore not suffer loss aversion.

10.5.1 Tax evasion.

Since Yitzahki, 1974, a number of researchers have pointed out the poor fit between a theory of behaviour based on expected utility on the one hand and the empirical regularities of tax evasion. In simple models of optimal evasion, taxpayers choose how much income to reveal to the authorities. If they fail to reveal all their income there is some probability of being caught and a consequential fine. Allingham and Sandmo, 1972, show that if p is the probability of being caught and the fine (as a proportion of the evaded income) is γ , then some evasion is optimal provided that $1 - p - p\gamma$ is positive.¹⁴ For many tax systems p is very small (2-3% according to Dhimi, and al-Nowaihi, 2006,) while γ is typically less than 2. Hence $1 - p - p\gamma$ is close to 1 and evasion is optimal. In fact, Andreoni et al, 1998, report that around 30% of (US) citizens evade taxes.¹⁵ In a similar vein expected utility theory underpredicts the level of tax evasion (i.e. how much tax is evaded) unless it is assumed that individuals are risk averse to a degree that does not match behaviour elsewhere. Thirdly, expected utility is a function of total wealth or income, levels of tax evasion should not be affected by whether or not taxpayers must pay some or all taxes in advance. In reality, paying tax in advance seems to reduce evasion (Chang and Schultz, 1990, using US. Internal Revenue Services data). Finally, if absolute risk aversion is declining in wealth, then higher tax rates should lead taxpayers to evade to a lesser degree. The opposite seems to be more realistic.

Bernasconi and Zenardi, 2004 and Dhimi and al-Nowaihi, 2006, apply versions of prospect theory to these paradoxes.¹⁶ In the latter, the authors demonstrate that cumulative prospect theory used with typical parameters, when allied to an assumption that the reference point is the post-tax, no evasion income and coupled with a stigma attached to evasion, can explain the four regularities. Leaving aside the introduction of a stigma term, the key assumption here is the choice of reference point. If the reference point was gross

¹⁴The smoothness of the utility function in income, means that the for infinitesimally small gambles the consumer is risk neutral. As a result, the consumer should evade some tax as long as the expected returns from doing so are positive. This yields the condition $1 - p - p\gamma$.

¹⁵One possible reason for this low percentage is that many salaried taxpayers have very little in the way of invisible income that could possibly be hidden.

¹⁶Using hypothetical payoffs in a sequence of evasion dilemmas, Schepanski, and Kelsey, 1990, find supporting evidence of an endowment effect in tax evasion.




Figure 10-7 Income-Leisure Trade-Off

income or post-tax with evasion, then their results would not follow. because all states of the world would be coded as losses. So, a critical but as yet unresolved question for policy framing is how reference points are established and maintained.

10.5.2 Labour Supply and Tax Framing.

Much of the debate on personal taxation over the last twenty years has been focused on its impact on incentives to work. We end this chapter with a theoretical discussion of the possible impact of framing effects on labour supply choices. As with much of the previous sections, the 'alternative' theory used to predict the possible impact of framing is very much guided by theories of reference dependent preferences.

Figure 10-7 shows a standard income-leisure trade-off diagram. The budget constraint has four important components:

- A minimum income level, b , which provides basic support even in the absence of paid employment;
- A terms of trade for converting leisure into consumption via paid work;
- A supplementary element of benefit, paid up to income y^l , above which
- There is a taper, at a rate of a , through which the supplementary element is gradually withdrawn, up to an upper level of income y^u .

This budget constraint mimics some aspects of tax/benefit systems such as those operating in the USA and UK, where working tax credits (WTCs) provide support for families and individuals with low incomes from the labour. In many cases the basic support for individuals out of work differs from the support level given to individuals in work, but for simplicity I have not made this distinction in the diagram.

Consider the four points, A-D, evaluated with regard to their suitability as potential reference points in the framing of policy. In doing so I restrict attention to frames which are truthful.¹⁷ The table below shows the view of each of them, as depicted as gains and losses from the other points. So, as viewed from B, point A is a gain in income, coupled with a loss in leisure. In the final column we also have a summary of the option in a positive language or spin.

What is clear from the discussion on reference point effects earlier in the book is that A appears to be the reference point maximizes the attractiveness of working, since all increases in leisure from A are accompanied by a change in consumption that is presented as a loss.

From	A	B	C	D	Positive Description
To					
A	-	Loss ($y^u - y^l$)	Loss ($y^u - b$)	Loss (y^u)	Work, income and independence from the State
	-	Gain ($l^B - l^A$)	Gain ($l^{CD} - l^A$)	Gain ($l^{CD} - l^A$)	
B	Gain ($y^u - y^l$)	-	Loss ($l^B - l^A$)	Loss (y^u)	Work, income and extra benefit
	Loss ($y^l - b$)	-	Gain ($l^{CD} - l^B$)	Gain ($l^{CD} - l^B$)	
C	Gain ($y^u - b$)	Loss ($l^{CD} - l^A$)	-	Loss (b)	Benefit Income b and maximum leisure (no work)
	Loss ($l^{CD} - l^B$)	Gain ($y^B - b$)	-	-	
D	Gain (y^u)	Gain (y^l)	Gain (b)	-	No income & maximum leisure (no work)
	Loss ($l^{CD} - l^A$)	Loss ($l^{CD} - l^B$)	-	-	

There are a number of important ways in which the actual policy context differs from this simple analysis of Table 10.2 and that make the choice of optimal reference point potentially complex. Chief amongst them is the fact that for many individuals more hours of work are not necessarily a bad thing, even leaving aside the additional income they can generate. In other words it is not straightforwardly true that increased leisure is the good. Work has often been supposed to bring psychological benefits and to be pleasurable for many individuals. Consequently, it is not unusual to see ‘U’ shaped indifference curves in labour economics texts.¹⁸ For the framer aiming to avoid the usual reductions in labour supply associated with income maintenance schemes this

¹⁷Some of the benefits of misleading boundedly rational individuals have been explored in Chapter Seven. Here, I am concerning with how the same information can be represented in a manner likely to change labour market behaviour. Such a constraint is potentially important because with the high effective marginal rates of taxation typical with working tax credit schemes all options involving higher labour supply can appear unattractive, no matter what frame is employed.

¹⁸Of course, in the absence of poverty and unemployment traps, the upward sloping section of the indifference curve will not typically be observable in the labour market. The point

has only benefits, since it means that, for instance, the choice of A versus D can be framed as a pair of gains ('you get to work more and you get income') rather than as a trade-off between one gain and one loss.

In addition to the multi-faceted nature of work another important dimension of the policy task is that typically wages will differ amongst the target population. To see why this makes a difference, consider one possible way of framing the consumer's choices:

'If you work 1 hour less, you will have £x less to spend each week.'

For an individual, for a given job, it is straightforward to calculate what x is. If he or she is on the section of the budget constraint between A and B, $x = (1-t)(1-a)w$, where w is the wage rate, t is the tax rate on labour market income and a is the 'taper rate' – the rate at which benefit is withdrawn for every post-tax £1 earned. When w varies across the population, then x will also vary. Now public documents and other messages framing benefit rules cannot typically be targeted to particular individuals, but rather must aim at some average potential beneficiary. This makes frames such as the one in the example either vague or irrelevant to a proportion of the target audience. Vagueness would tend to undermine the credibility of the framing effect and would therefore be more likely to leave the individual's prior frame intact.

Some frames may be less affected by the restriction that the message be suitable to a wide range of individuals. For instance 'if you work more than 16 hours you may be entitled to up to £100 of extra help each week', does not require mention of a specific wage rate to be truthful. However, it does not say anything about the taper. To be wage independent, that statement could be of the form:

'as you work more and earn more, the support you receive will be withdrawn, at a rate of £a for every £1 extra you receive in your pay packet.'

On the other hand, its converse: if you work less, then for every £1 less you receive in your pay packet you will receive an extra £a in support, sounds far more complicated and therefore less convincing as a frame.

In both these last two examples, independence from mentioning the wage rate is achieved by separating out the benefit and the post-tax income in the description of the choice set. Doing so opens up the possibility of manipulating the set of mental accounts as well as the perception of choices as gains or losses. We have seen how mental accounts can alter choices. It appeared that simple rules such as separate out gains, integrate small losses into large gains, integrate losses and (possibly) separate out large losses and small gains could alter the perceived utility to be had from choices. Here, the receipt of benefit is potentially separable from the receipt of income from work. If we were to follow the simple prospect theory rules, then to maximize the favourability of A over C or D would require either integrating the loss of benefit associated with a gain in wage income (if C or D is the reference point), and also, possibly, integrating the gain in benefit associated with a loss of wage income when A is the reference point.

is that in drawing preferences in such a manner, the authors are recognising the inherently dual nature of the psychology of work.

More could be said on this issue. In the UK for instance, the Labour government is separating out the support given to members of working but low income families into the basic working tax credit (WTC) available to all low income households supplying more than 4 hours a week of paid labour and a child tax credit (CTC) paid as a supplement to wage income for low income families with one or more children. Consider the following hypothetical example, where an individual in a household with children must choose between the status quo and working one extra hour per week.

	Work	Receive	Lose	Lose
Old	1 extra hour	£5 extra wage income	£2.75 family tax credit	-
New	1 extra hour	£5 extra wage income	£1.85 tax credit	£0.90 child tax credit

Financially these two presentations are equivalent. In both of them the subject has £2.25 available from working one extra hour. On the other hand, what was once one gain of income and one loss of income is now one gain of income and two losses. If the choice is framed in this way, then prospect theory would predict that working longer hours would be less attractive under the new frame compared with the old. On top of that, some of the lost income is specifically linked to children, which may evoke a quite different sense of loss.

Framing effects may also impact on the decision to work. Consider the following two options, A and B.

- Option A: Unemployed; current unemployment benefit of £x.
- Option B: Full-time Employed; a wage of £y

Suppose someone was asked, which would you choose, option A or option B? Compare the result to the question C:

- C: Suppose you are currently unemployed and receive £x of unemployment benefit per week. You are offered full-time employment. What is the minimum wage you would be willing to accept for that employment?

The evidence on preference reversal (e.g. see Cubitt et al, 2004), would suggest that many individuals would give contradictory answers to these two types of questions. In particular, if the employed/unemployed dimension of the problem is viewed as more prominent than the income dimension, then there would be a tendency for individuals to set a reservation wage in their answer to question C which would be higher than values of y for which they would accept option B. The evidence on loss aversion would also suggest a tendency for individuals to choose option A when endowed with option A and to choose option B when endowed with option B. In both cases, therefore behaviour in the labour market would be sensitive to the framing of the choices available to the individual concerned.

There is one final point to be emphasised. It was noted above that marginal tax rates can be very close to 100% with tax credit schemes. In some cases rates exceed 90% (in the UK for instance - see Connolly and Munro, 1999). Consider two frames for the individual facing a marginal rate of 95% for instance, with an hourly wage rate of £5:

- If you work one hour more you will gain an £0.25 extra income
- If you work one hour less you will lose £0.25.

It is not obvious that either frame presents the choice between fewer and more hours in more positive light. Consequently, it is worth reminding ourselves that even the most sophisticated of framing strategies may have little impact on behaviour if one choice is so clearly better than another.

10.6 Conclusion.

In this chapter I have explored some non-standard fiscal policies. Government often seem to use them unwittingly (e.g. the Winter Fuel allowance) and there has been little in the way of an analytical approach to their employment within economics. The apparent success of Thaler and Benartzi's savings scheme does suggest that insight from the behavioural and experimental economics literature can be used to design and understand framing effects.

I have also applied some of the ideas from the early sections of the book to the issues of pensions and savings. In some ways it is easier to do this with behaviour in financial markets than in other contexts. Partly this is because anomalies are well-established in financial markets and this in turn is due to the fact that the standards of rationality are somewhat clearer in situations where risk and return are the prime factors determining choice in the rational choice model.

As in other chapters, the emphasis has been very much on the individual and the framing of individual choice, but it is worth ending with the reminder that much behaviour takes place in a social setting and that the government often has the ability to frame the social context of decisions as well as the individual frame. The positive impact on recycling of using transparent bags for household rubbish is predicated on the notion that individuals will wish to avoid shame of public exposure. Other deterrent policies, from tax evasion to littering often seek to exploit the same sorts of emotions. There is therefore much more to be learnt on the effects of policy framing.