13

The Taxation of Household Savings

The taxation of savings plays a central role in how economists evaluate a tax system. There are five reasons for this.

First, the way in which savings are taxed is a key characteristic of the *tax base*. If the tax base is defined as including income from savings as well as labour earnings, and all components of the tax base are taxed equally, this yields the so-called 'comprehensive income tax'. Alternatively, if earnings that are saved, and the returns to savings, are not taxed until such time as they are used for consumption, then the resulting tax system will be an 'expenditure tax' or a 'consumption tax'. The difference in the tax treatment of savings is the critical difference between these two tax bases.

Second, the tax treatment of savings is an important determinant of the extent to which the tax system recognizes interpersonal differences in *lifetime income*, as opposed to annual income. Careful design of the tax treatment of savings is one way of trying to equalize the tax burden on taxpayers with similar lifetime incomes but different income patterns over their life courses.

Third, taxation of savings stands right on the boundary between taxation of personal income and taxation of company profits. How we tax savings can influence the behaviour of small firms and the self-employed, as well as the allocation of capital to large firms.

Fourth, savings taxation can affect both the total amount of savings in the economy and, probably more importantly, how those savings are allocated across different assets. This can directly affect the amount of capital invested and how efficiently it is invested.

Finally, for individuals, the taxation of savings affects their decisions on how much to save, when to save, and how much risk to take when allocating their savings between assets. It therefore directly affects their welfare and particularly their welfare in periods of retirement or unemployment, when they may need to rely on accumulated savings.

This chapter examines the economic arguments for different possible systems of savings taxation. In the next chapter, we focus on some specific directions for reform in the UK context.

We start our discussion with some evidence on people's actual saving behaviour. This is important in understanding what we might want to achieve through taxation. We go on to look at the case for exempting tax on the normal return to savings. The case for exempting the normal return to savings from taxation is likely to depend on, among other factors, the reasons that people save in the first place. Many do so in order to consume at one period of their lifetime rather than another. By sacrificing consumption today, saving is a way of generating future income and, like other forms of investment, there is a case for exempting the normal return. The taxation of the normal return to savings distorts the timing of lifetime consumption and labour supply. A timing-neutral tax system would not create such distortions, and there are a number of tax systems that achieve such neutrality.

A consumption tax does not create distortions in the timing of consumption, while a comprehensive income tax does. This is because the latter reduces the after-tax rate of return relative to the pre-tax return, and because the rate of return the consumer receives determines the effective price of future versus current consumption. Since one of our objectives is to avoid distorting intertemporal choices, at least for a large fraction of the population, we explore three possible routes to savings taxation that maintain neutrality over when consumption occurs. We explain these and contrast them with some of the difficulties inherent in a comprehensive income tax and with the additional distortions—for example, over whether returns to savings are taken as income or capital gains—that are almost inevitably introduced by such a tax.

In spite of a vast body of research on the appropriate taxation of savings, we recognize at the outset that economic theory does not provide an unequivocal recommendation on the issue of optimal tax design. We

therefore rely in part on broadly-attractive concepts, such as tax neutrality, in framing our analysis. We view neutrality as a constructive benchmark in understanding the issues surrounding the design of savings taxation. There is potentially a rich array of ways in which individuals differ with regard to saving behaviour based on underlying preferences and opportunities. In the absence of such detailed knowledge, it seems sensible to begin from this benchmark and look for justifications for deviating from it.

The different routes to neutral taxation involve collecting taxes at different times. In simple terms, one route involves collecting tax up front and not taxing the later return to savings. Another route involves not levying tax on any income that is saved, but then taxing withdrawals (rather as pensions are taxed in the UK today). A third route is to exempt a 'normal' return to savings but to tax 'excess' ('supernormal') returns. These obviously have different cash-flow consequences for governments.

An important further difference between these systems from the point of view of individuals arises when income taxes are progressive or, in general, when individuals expect to face different marginal tax rates at different times of their life. Then the different systems will have different effects on people's incentives to save according to the pattern of their income and consumption over time. One possibility we examine is to allow people to choose between the different systems and thereby to smooth their taxable income between periods. In some circumstances, that can move us towards the ideal of taxing lifetime income.

When considering the taxation of savings, it is also important to consider the taxation of borrowing (negative savings) and the taxation of 'human capital'. If financial investments and investments in the future through education are treated differently, then choices may be distorted across this margin.

Having discussed tax structures that achieve tax neutrality, we conclude this chapter by examining the economic case for tax neutrality. We note that there are a number of potential justifications for deviating from the timing-neutral tax benchmark, and observe that the optimal taxation of capital income is a very active area of ongoing research.[1] A number of recent studies suggest that the economic case for taxing the normal return may be more

[1] See Banks and Diamond (2010) for a review.

ambiguous than many analysts have suggested, and, as a result, our conclusions on this issue cannot be completely clear-cut. We do retain neutrality as a useful benchmark and suggest that there are many practical reasons for assigning a presumption to a neutral system.

In the next chapter, we go on to apply some of these insights to the current UK policy context, with some recommendations for changes to the tax regime as it applies to particular asset classes. In both these chapters, we focus only on life-cycle savings—that is, savings that are accumulated in one part of an individual's life in order to increase consumption at a later date. The analysis does not necessarily follow through to situations in which savings derive from, or are used to provide, gifts and inheritances. Motives for bequests and the extent to which individuals save in pursuit of dynastic wealth are poorly understood. We look at the taxation of wealth transfers separately in Chapter 15.

13.1. SAVING BEHAVIOUR

It is worth starting by looking at some general evidence on people's actual saving behaviour. If we were to find that people neither make any attempt to smooth their consumption over their lifetime, nor change their behaviour in the face of different taxation of different assets, we might conclude that how savings are taxed matters little. What we show here is that, in fact, people generally do both these things—suggesting that taxation does matter. There is also a very extensive formal literature that confirms the impressionistic evidence presented here.[2]

13.1.1. Saving over the Life Cycle

In broad terms, people tend to save less when their incomes are low and their needs are high—for example, when they have children and take paternity or maternity leave—and save more when incomes are high and needs are low—the period between children leaving home and retirement, for example. On

[2] See Poterba (2002), Attanasio and Wakefield (2010), and Attanasio and Weber (2010).

the whole, many people do a fair job of maintaining stable consumption levels during their working life and in retirement. Of course, this is not true for all. Government policy in general, and tax policy in particular, cannot rely on individuals always making optimal saving decisions. One consequence is that a balance has to be struck between avoiding distortions to saving behaviour and providing a safety net for those who do not prepare well for the future.

Individuals save (or repay debts) or dissave (either borrowing or running down their existing wealth) when the amount they choose to consume differs from the amount of income they receive in a particular time period. As Figure 13.1 shows, people's incomes tend to rise and then fall on average over the course of their lives. Although this pattern also holds for consumption (here measured by expenditure on non-durables and services), there is much less variation. Consumption of non-durables and services is flatter than income over the life cycle.

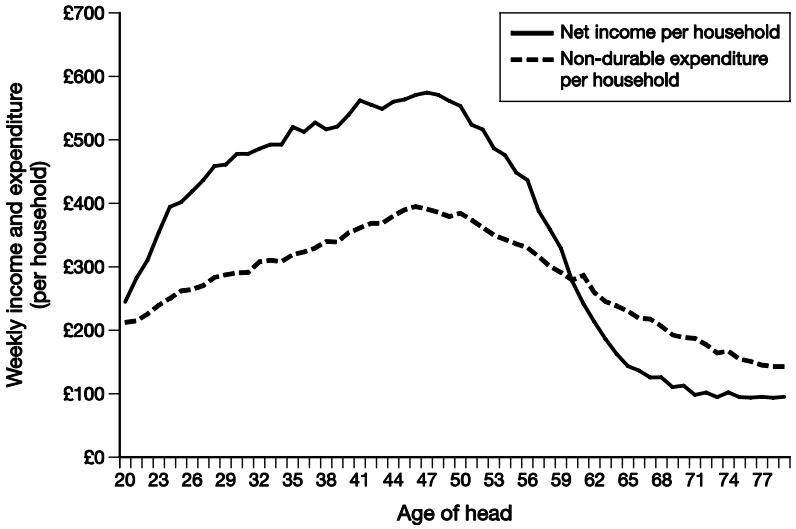

Figure 13.1. Net income and net expenditure per household

Notes: Average weekly net income is after-tax-and-National-Insurance take-home pay plus benefits and other unearned income. Average weekly expenditure on non-durables and services.

Source: Authors' calculations using the IFS tax and benefit microsimulation model, TAXBEN, run on uprated data from the UK Family Expenditure Survey / Expenditure and Food Survey, 1974–2008.

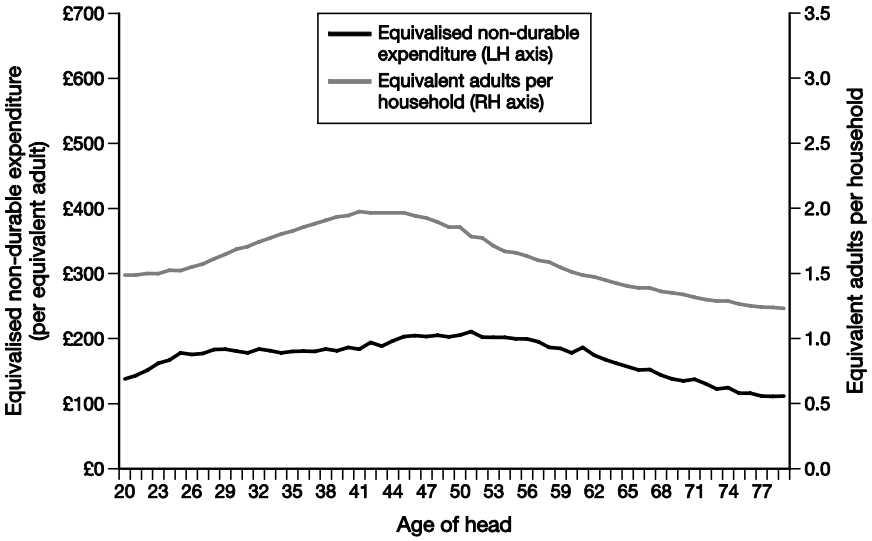

Figure 13.2. Consumption and needs

Notes: Average weekly expenditure on non-durables and services. Number of equivalent adults is computed using the modified OECD scale.

Source: Authors' calculations using the IFS tax and benefit microsimulation model, TAXBEN, run on uprated data from the UK Family Expenditure Survey / Expenditure and Food Survey, 1974–2008.

Of course, a family will not aim to smooth consumption exactly; it will wish to vary consumption with family size. Figure 13.2 shows how the number of 'equivalent adults' in a typical household varies over the life cycle. This is a measure of household size that takes into account the fact that a child is less costly to support than an adult.[3] The figure also shows how expenditure per equivalent adult varies much less than family size and is reasonably constant across families with heads of different ages. This is direct evidence of the way that consumption is smoothed to adapt to needs over time.

The net effect is that there is usually a desire to borrow, then to save, and then to draw down savings as the stages of the life cycle progress. Figures 13.1 and 13.2 suggest that people on average do a reasonably good job of

[3] An equivalence scale assigns weights to households of different compositions, intended to reflect the different resources they require to reach the same standard of living. We use the modified OECD scale in which the first adult in a household has a weight of 1 and the second and subsequent adults each have a weight of 0.5. Children aged 14 and over also have a weight of 0.5, and children 13 and younger have a weight of 0.3.

'consumption-smoothing', once we adjust for changes in needs as family size changes. But, of course, people's behaviour is not quite as straightforward as this account suggests.

The standard economic model assumes that consumers make sophisticated decisions based on well-founded expectations and beliefs about future economic events. In reality, even if people think hard about the long-term decisions they face, they are likely to take decisions on the basis of only limited information.

Some individuals and families will find consumption-smoothing hard to achieve, especially if they have limited access to credit (which is more often the case for younger and poorer households than for older and richer ones). Others may be myopic, so they save too little for the future and have to either consume less in retirement or delay retirement. More specifically, people's decision-making may be driven less by long-term thinking and more by the desire for immediate gratification than the traditional model assumes.[4]

It is, perhaps, not surprising that apparently myopic behaviour occurs most often among individuals and families with relatively poor educational qualifications and low wealth. Recent experimental studies suggest that individuals with higher ability (and earning potential) tend to be more patient and better able to make complex decisions.[5] Intellectual ability and numeracy are both associated with higher likelihoods of holding stocks and of having a private pension, and not just because able and numerate people tend to have more financial wealth in total.[6]

There is particular policy concern over the extent to which people save enough to support themselves, and maintain their standard of living, in retirement. There is, in fact, a well-documented fall in consumption at the time of retirement.[7] While two-thirds of this drop can be explained within the context of a life-cycle consumption plan (for example, a fall in work-related expenditures, or less spending on expensive prepared foods as people have more time to cook), the remaining third does appear to indicate that some people do not save enough for retirement. Concerns about

[4] See Ainslie (1975) and Thaler and Shefrin (1981).

[5] See e.g. Parker and Fischhoff (2005), Kirby, Winston, and Santiesteban (2005), and Bettinger and Slonim (2006).

[6] Banks and Oldfield, 2006.

[7] Banks, Blundell, and Tanner, 1998.

undersaving for retirement have led the UK government to propose to 'nudge' people into saving by ensuring that everyone is automatically enrolled into an employer-sponsored pension. They will have to make an active decision not to save, rather than an active decision to save.

In designing the taxation of savings, we need to recognize that patience, self-control, and the ability to take long-term decisions in a sensible way vary from person to person. We cannot rely on all individuals to make considered provision for their long-term needs through their own private decisions.

It is this apparent lack of rationality by some people that can drive government policy on savings, pensions, and social insurance. At one extreme, government can simply tax everyone in work and provide incomes in retirement that are unrelated to tax payments. But linking benefits to contributions can improve the efficiency of a tax system.[8] At the other extreme, where future provision is provided solely through private savings and private insurance contributions, there would no longer be any distortions to the timing of consumption produced by such contributions. However, there are limits to the ability of individuals and families to make life-cycle provisions through voluntary insurance contributions and private savings. The income tax and benefit system will continue to be called upon to provide a floor for living standards.

13.1.2. Allocation of Savings between Assets

We need to distinguish between the issue of how much people choose to save and the variety of assets or financial instruments through which they save. The most recent estimates from the Office for National Statistics (ONS)[9] are that households in the UK held around £9 trillion of wealth in 2006–08, of which 39% was held in private pensions and a similar amount in property, largely owner-occupied housing. Pensions and housing are, of course, two relatively tax-favoured ways of holding wealth.

Evidence from the English Longitudinal Study of Ageing (ELSA) provides more detail on the distribution and composition of savings, specifically for

[8] See Bovenberg and Sørensen (2004) and Bovenberg, Hansen, and Sørensen (2008).

[9] Office for National Statistics, 2009.

Table 13.1. Fraction of financial wealth held in different assets in England, 52- to 64-year-olds, 2004

Decile of gross financial wealth	Range of gross financial wealth (£'000s)	Percentage of wealth held in:		
		Private pensions	ISAs, PEPs, and TESSAs	Other assets
Poorest	<1.7	12.6	9.1	78.3
2	1.7–16.6	54.8	13.8	31.5
3	16.6–39.1	65.2	11.0	23.8
4	39.1–75.9	68.2	10.8	21.0
5	75.9–122.3	69.7	7.9	22.3
6	122.3–177.2	74.7	6.8	18.5
7	177.2–245.4	78.1	6.2	15.7
8	245.4–350.3	81.8	4.6	13.6
9	350.3–511.2	79.0	5.7	15.3
Richest	>511.2	68.4	4.4	27.3
All		73.6	5.5	20.9

Notes: Benefit units with at least one member aged 52–64 in the 2004 English Longitudinal Study of Ageing. Private pension wealth comprises current fund value of defined contribution (DC) pensions and the value of accrued entitlements to date of private defined benefit (DB) pensions (based on assumption of no further real earnings growth). Percentages are ratios of means for each decile group, not group means of individual ratios. Numbers do not always sum exactly, due to rounding.

Source: Wakefield, 2009.

individuals aged over 50.[10] Table 13.1 considers the non-housing wealth of individuals in a key part of their lifetime as far as savings are concerned—those aged 52–64. Those with low levels of gross financial wealth hold the largest proportion in 'other assets', mainly interest-bearing accounts, which are subject to taxation as interest is credited and have no tax advantage when savings are deposited. These are the most heavily taxed savings assets, particularly in an inflationary environment.

When the tax system changes, people respond. When tax-privileged accounts such as Individual Savings Accounts (ISAs; and their predecessors TESSAs and PEPs) were introduced, billions of pounds rapidly flowed into them.[11] In the late 1980s, the government introduced a reform that provided

[10] Banks, Emmerson, and Tetlow, 2007.

[11] Attanasio and Wakefield (2010) and references therein.

young people with very large incentives to 'contract out' of the State Earnings-Related Pension Scheme (SERPS) into personal pensions. Young people did exactly that, with 40% of those in their 20s moving into personal pensions, along with remarkable numbers of 16- to 19-year-olds. As the incentives were withdrawn, so coverage fell. In the early 2000s, there were also marked responses to changes in tax limits that accompanied the introduction of stakeholder pensions.[12]

The key point is not that everybody at all times responds rationally to tax incentives, but rather that there is compelling evidence that such incentives are major drivers of individual decision-making and of the allocation of resources across the economy. Large and salient changes in the savings tax system change behaviour.

13.2. THE NEUTRALITY PRINCIPLE

As the discussion of actual saving behaviour has illustrated, two distinct concepts of neutrality matter with respect to the taxation of savings. The first is neutrality over the level and timing of saving—if the tax system is neutral in this sense, it does not distort people's choice over when to consume their income. The second is neutrality between different types of savings vehicles or assets—a neutral tax system in this sense does not distort people's choices over the assets in which they save.

A tax system that levies a tax on the 'normal return to savings'—the return that just compensates for delay in consumption (without any additional return related to risk-taking, for example, which we discuss below)—cannot satisfy the first neutrality criterion. It taxes people who choose to consume later in life more heavily than people who choose to consume earlier in life.

Taxing the normal return to savings means taxing consumption tomorrow more heavily than consumption today. In some contexts, having different tax rates on consumption according to when it occurs is conceptually rather like having different tax rates on different forms of consumption. The arguments over the merits of such different tax rates would then be directly

[12] Chung et al., 2008.

parallel to the arguments discussed in Chapter 6 over whether to tax some consumption goods more heavily than others.[13]

Recall that, in an ideal world, we would like to tax people according to their lifetime earning capacity—broadly equivalent to their potential consumption. The problem for policymakers is that ability cannot be observed directly, so we use actual earnings or expenditure as an imperfect proxy, which has the unfortunate consequence of discouraging people from earning (or spending) as much as they would otherwise like: we distort their behaviour towards choosing more leisure time instead. Taxing the normal return to savings can only improve on this if it allows us to target high-ability people more accurately or with less distortion to labour supply.

It might appear that taxing savings is an effective way to redistribute—after all, aren't people with large savings wealthy almost by definition? But someone with savings is not necessarily better off over their lifetime than another person without savings. The two might earn and spend similar amounts over their lifetimes, but at different times: one earns his money when young and saves it to spend when he is old, while for the other the timings of earning and spending are close together. We can tax people on their total resources by taxing their money at its source (taxing earnings) or when it is finally used for consumption (taxing expenditure).[14] We can tax better-off people more heavily by making the rate schedule applied to earnings or expenditure more progressive. If—given what we already know from their actual income and expenditure—people's saving decisions tell us nothing more about their underlying earning capacity, just about their taste for consuming tomorrow rather than today, then taxing savings cannot help us to target high-ability people more accurately than taxing earnings or expenditure. By taxing the normal return to savings, we are not taxing the better-off; we are taxing those who spend their money tomorrow rather than today. That seems both unfair and inefficient, unless there is a relationship between when individuals choose to spend their money and other attributes that might be a basis for taxation, such as their underlying earning capacity.

[13] See Atkinson and Stiglitz (1976).

[14] If people inherit money rather than earn it, or bequeath it rather than spend it, then (although ultimately the money must have been earned by someone and must be spent by someone) different considerations apply. These are the subject of Chapter 15.

Broadening the tax base to include savings might seem like it allows us to reduce tax rates on earnings and reduce disincentives to work. But work decisions involve trading off consumption against leisure. If someone is working in order to finance *future* consumption, then taxing savings—reducing the future consumption that can be bought with earnings—discourages work just like taxing earnings directly. Why discourage work more among those who prefer to consume the proceeds later?

Arguments about consumption today versus tomorrow only apply to taxation of the normal return to savings—the return that just compensates for delaying consumption. In Section 13.2.2, we will see that there are strong arguments for taxing returns in excess of this. In Section 13.3, we will discuss cases where the logic even for exempting the normal return breaks down—for example, where people's saving decisions do tell us about their earning capacity, or where taxing future consumption does not reduce labour supply in the same way as taxing current consumption—and consider whether such cases justify departing from neutrality in practice. But neutrality over the timing of consumption is, at the very least, a reasonable starting point for tax design.

The second type of neutrality—neutrality between different assets—is lost if different assets (housing, pensions, other financial assets) are taxed differently. One would generally need rather strong reasons for deviating from this form of neutrality—tax policy shouldn't really be influencing whether I decide to save in gilts, shares, or a savings account. One potentially substantial exception is that there may be a case for treating pensions more generously than other forms of savings so as to encourage people who may not plan well for the long term to save for retirement in a form that will provide them with a regular income.

While there are limitations to the standard arguments in support of both neutrality concepts, and particularly the first, understanding what types of tax system will achieve neutrality is a natural benchmark for any reform discussion.

We look now at why a comprehensive income tax cannot achieve either of these types of neutrality, then outline three different approaches to achieving neutrality, before going on to look at complications to this story created by income tax systems with more than one tax rate.

13.2.1. Why Standard Income Taxation Cannot Achieve Neutrality

A standard income tax treatment of savings achieves neutrality neither over time nor across assets. An income tax deters saving by making future consumption more expensive than current consumption. Because it taxes earnings as they are received and then taxes any return to savings, the present value of the income is greater if it is used for consumption now than if it is used for consumption in the future. Furthermore, unless there is full indexation for inflation, the degree to which this occurs will vary over time in an arbitrary way with fluctuations in the rate of inflation because the nominal return will be taxed, not just the real return. If inflation is high, interest rates will tend to be high in nominal terms to compensate for the fact that the real value of the principal will be falling. Taxing that nominal return implies that the effective tax rate on the real return to interest-bearing assets tends to increase with the rate of inflation.[15]

The phenomenon of compound interest means that a tax that reduces the effective rate of return on savings looks increasingly penal—reduces the final wealth generated more—the longer the time horizon involved. For a young person saving for much later in life, this can make a startling difference to the value of wealth generated by a given amount saved. Even ignoring inflation, a tax that reduces the net interest rate on a bank account from 5% to 4% will reduce the value of the account by around 1% after one year (from £105 to £104), but by around 9% after ten years and by 38% after 50 years. For quite plausible saving over an individual's life, the combination of inflation and compound interest means that standard income taxes reduce the future consumption that can be bought by forgoing consumption today to a far greater degree than one might suspect from looking at statutory tax rates.

It is difficult to design an income tax that is neutral across assets, particularly when capital gains are taxed at realization and without any adjustment that makes a realization-based tax equivalent to a tax on accruing gains. Capital gains are a return to savings in just the same way as interest

[15] It is possible to design a tax system based on realization accounting that achieves a uniform capital tax. Indeed, Auerbach and Bradford (2004) develop a generalized cash-flow tax that avoids having to measure capital income while at the same time effectively imposing an income tax at a constant rate on all capital income.

income or dividends. Under a comprehensive income tax, capital income (including capital gains) should be taxed as it accrues, or in a way that is equivalent to accrual taxation. So capital gains need to be taxed at the same rate as other components of income, which is clearly possible (though not what happens in the UK). In the standard formulation of a comprehensive income tax, capital gains are taxed at the same time as other forms of income from savings. That implies taxation on accrual (when the rise in value occurs) rather than on realization (on disposal of the asset). For an asset that increases in value and is then held for several more years before being sold, the effect of taxation on realization is to defer the tax payment on the accrued capital gain for several years. While it is possible to design realization-based capital gains taxes that provide investors with the same incentives as a tax on accruals, such taxes would require modifying the asset's tax basis by an amount that depends on rates of return since the asset was purchased.

Deferring or delaying tax payments is valuable to taxpayers—this can be thought of as the equivalent of an interest-free loan from the government to the taxpayer, from the time the asset increases in value to the time it is sold. This delay reduces the effective tax rate on capital gains, particularly for assets that are held for long periods. This unequal treatment favours assets that generate returns in the form of capital gains over assets that generate returns in the form of cash income. This also creates incentives for cash income to be converted into capital gains, which may be particularly important in the context of business assets, and therefore favours some individuals over others. Taxing capital gains on realization without any 'accrual-equivalent' adjustment also creates a 'lock-in' effect—once an asset has risen in value, there is an incentive to hold on to it, to shield the accrued gain from tax for a longer period. Taxing capital gains on accrual would, though, be extremely difficult for two reasons: first, all assets would need to be 'marked-to-market' or valued in periods when they are not traded; and second, individuals may be required to pay tax on accrued gains in periods when they lack the liquid financial resources to make these payments.

In practice, then, taxing the return to savings under a standard income tax implies accepting arbitrary distortions to the pattern of saving both over time and across assets. As we shall see, an expenditure tax avoids distorting the choice between assets that yield cash income or capital gains, and the

holding-period distortion, even though gains are taxed only when they are realized and consumed.

13.2.2. Alternative Routes to Savings-Neutral Taxation

A comprehensive income tax cannot take us to a savings-neutral system of taxation. But there is in fact more than one route to a savings-neutral system. We consider three here. In doing so, and in order to facilitate the discussion, we find it very useful to make use of some simple notation. We describe each stage in the life of the asset in which savings are invested as taxed (T) or exempt from tax (E).

There are three stages to consider: first, when income is received (i.e. before or at the point at which it is paid into a savings account or used to purchase an asset); second, as the returns (interest, capital gains, or distributable profit) accrue; and third, when the funds are withdrawn from an account or an asset is sold.

In this notation, a (cash-flow) expenditure tax is defined as EET. Tax is simply paid on income used for consumption at the time the expenditure is made. This is equivalent to saving in a tax-deferred account and most pension saving operates in this way. In contrast, the comprehensive income tax is TTE. That is, savings are made out of taxed income; all returns are taxed, including the normal return; but no further tax is due when the savings are withdrawn.

With these concepts in mind, there are three potential alternative savings-neutral forms of taxation. They are:

- a 'cash-flow expenditure tax', which taxes only income used for consumption when it is spent—EET;
- a 'labour earnings tax', which excludes all savings income from taxation, but with no exemption for savings when first made—TEE;
- an 'income tax with a rate-of-return allowance' (RRA), which taxes labour earnings and supernormal returns to savings—TtE. The lower-case 't' in TtE denotes the exemption for the normal return.

The three savings-neutral approaches are broadly equivalent in the absence of supernormal returns. All three leave the normal risk-free return untaxed

and consequently leave the choice between consumption today and consumption tomorrow undistorted.[16]

The different forms do, however, have different implications for the tax treatment of returns in excess of the normal return, as well as for the time path of government revenue. The normal return is a central concept here. It can be thought of as the return obtained by holding savings in the form of a safe, interest-bearing asset. For this reason, it is often called the normal risk-free return.[17] It is this return that we want to avoid taxing in order to avoid distorting decisions over the timing of consumption. It is because it taxes the normal return that the income tax distorts these decisions.

Returns above the normal return may reflect differential risk across different investments or some form of rent earned by investors. The source of excess returns may have an important effect on the economic consequences of different approaches to achieving neutrality. The earnings tax (TEE) leaves excess returns untouched by the tax system. It doesn't matter how well my investments do, I pay no further tax. The expenditure tax (EET) and rate-of-return allowance (TtE) bring excess returns into the tax base (and both raise revenue by taxing rents). This is a crucial difference. Widespread application of the TEE system would allow successful investors to earn unlimited rewards without being taxed. It is quite inappropriate as a general regime for business assets and other risky investments. The TEE regime, of course, also requires a very sharp differentiation between earned income and investment income, since the former is taxed and the latter not.

Earnings tax (TEE) and expenditure tax (EET) treatments of savings are widely used for certain assets. Private pension plans in the UK approximate an EET treatment. This is also the case for human capital investments where the investment of time in education is not taxed but the returns are. Roth 401(k) plans in the US and ISAs in the UK are examples of assets that are given a TEE tax treatment. Owner-occupied housing in the UK and most durable consumption goods attract a TEE treatment too, since they are bought out of after-tax income but there is no tax paid on returns, even 'excess' returns.

[16] At least for consumers who can borrow at the normal return and face a constant tax rate over time (we address these caveats in Section 13.2.3).

[17] In most developed countries and most time periods, this can be well approximated by the interest rate on medium-maturity government bonds (Sørensen, 2007).

A standard income tax (TTE) taxes all the returns from capital investments, including the normal return.

An EET base can be thought of as giving tax relief for saving up front. The rate-of-return allowance can be viewed as an expenditure tax with deferred rather than immediate tax relief for saving. Their common feature is that, unlike the comprehensive income tax (TTE), they do not tax the normal return to savings. Indeed, the RRA and the EET can be viewed as two special cases of a more general savings-neutral tax base.[18,19] The RRA has gained increasing attention in the economic literature and has been introduced in Norway.[20] We believe it should be taken seriously in the savings tax reform debate. It achieves the neutrality we are looking for. It has the potential to be less disruptive to implement than a traditional consumption tax. It maintains government revenues up front. And it allows the same tax rates to apply to (above-normal) returns to savings as apply to labour income.[21]

These different tax regimes for savings can all be applied to borrowing as well, as described in Box 13.1.

To help understand the different systems, we develop a simple example that compares a standard income tax (TTE) with the three alternative savings-neutral tax regimes. In our example, we look at an individual who is considering saving in an asset that provides a 5% annual return. For every £100,000 of this year's income saved, the following year there is interest income of £5,000 plus principal of £100,000, a total of £105,000.

A standard income tax at 20% gives tax on interest income of £1,000, after-tax interest income of £4,000, and a return of only 4%. This is a disincentive to save. The TTE case is displayed in the first column of numbers in Table 13.2. In the remaining columns, we draw out the comparisons for the savings-neutral tax systems.

[18] In much the same way that cash-flow corporation taxes and ACE-type taxes are two special cases of a more general investment-neutral corporate tax base (see Chapter 17).

[19] An intermediate case would give immediate tax relief for part of the individual's net saving, with deferred tax relief (with the same present value) for the remainder.

[20] Sørensen, 2009.

[21] It should be added that the full 'general equilibrium' effects of moving between these different savings-neutral tax systems still need to be fully worked out.

Box 13.1. Tax regimes for borrowing

Borrowing can be thought of as negative saving, and the same four tax treatments we consider for savings could all, in principle, be applied.

- **TEE—a labour earnings tax** simply ignores borrowing, like it ignores saving. Neither taking out a loan, nor making payments of interest or principal, has any effect on tax liability.
- **EET—a cash-flow expenditure tax** involves taxing all cash inflows and deducting all outflows, hence adding the loan to taxable income for the year when it is taken out and then deducting all payments of interest and principal.
- **TTE—a comprehensive income tax** treatment of borrowing allows full deductibility of interest payments from taxable income (but does not add the amount borrowed to taxable income or deduct repayments of principal), just as it fully taxes interest income on savings. A comprehensive income tax thus taxes saving and subsidizes borrowing.
- **TtE—a rate-of-return allowance** regime allows deductibility of interest payments, like TTE, but only in so far as they exceed a 'normal' rate of interest on the outstanding principal. (Unlike with TTE, there is no difference in present-value terms between making interest payments and making repayments of principal. If a payment is labelled interest, it is deductible; if it is labelled principal, it is not deductible but, by reducing the value of the outstanding debt, it reduces the stream of 'normal' interest allowances to offset against future interest deductions.)

Table 13.2. Comparison of savings tax regimes with normal returns (assumed 5%)

	TTE	TEE	EET	TtE
Purchase price	100,000	100,000	100,000	100,000
Tax relief in year 1	0	0	20,000	0
After-tax contribution	100,000	100,000	80,000	100,000
Value of asset in year 2	104,000	105,000	105,000	105,000
After-tax withdrawal	104,000	105,000	84,000	105,000
Tax paid in year 2	1,000	0	21,000	0
Present value of year 1 tax relief	0	0	21,000	0
Present value of tax paid	1,000	0	0	0

Under an earnings tax (TEE), the purchase again costs £100,000 in terms of consumption forgone, but no tax is then levied on the return, so £105,000 can be withdrawn. An expenditure tax (EET) can be thought of as providing a tax relief of 20% on the purchase price. Hence the cost of the asset in terms of consumption forgone is £80,000. That is, the expenditure tax gives tax relief of £20,000 on saving of £100,000 in the first year. It then taxes the withdrawal of £105,000 in the second year, resulting in a tax payment of £21,000. After tax, the saver gives up £80,000 this year and gets £84,000 next year, a return of 5%. Put another way, the present values of tax relief in period 1 and tax payment in period 2 are equal. There is no distortion to the intertemporal allocation of consumption.

Now suppose that instead of giving tax relief of £20,000 this year, we carry this forward, marked up at the interest rate of 5%, and give tax relief of £21,000 next year. The saver then gives up £100,000 this year and gets £105,000 next year, just as in the TEE case, a return of 5%. This is displayed in the final column of Table 13.2. The EET and TtE approaches are equivalent provided the individual is indifferent between tax relief of £20,000 in year 1 and tax relief of £21,000 in year 2. We can achieve this here, and more generally, by providing a rate-of-return allowance, calculated as the risk-free (nominal) interest rate multiplied by the stock of savings (at historic cost) at the end of the previous year—5% of £100,000 = £5,000 in the example.

The situation changes when there is a return above the normal rate. To illustrate, suppose that the normal return is 5% but that the asset purchased provides a return of 10%. We assume in this case that the excess return is a rent earned by the investor. This situation is illustrated in Table 13.3 for each system that we are considering. We see that TtE and EET are equivalent, while TEE is different. This is because under TtE there is only an allowance of £5,000 in year 2 to set against the return of £10,000. In this case of supernormal returns, the return above £5,000 is taxed at 20%.

This stylized example is useful for understanding basic principles, though of course there are other important differences between the systems. For example, in the case of a risky asset, both the timing and riskiness of government revenue receipts are different between the systems. With the TEE treatment, all revenues are certain and are received in the first period.

Table 13.3. Comparison of savings tax regimes with excess returns (assumed 10% with normal at 5%)

	TTE	TEE	EET	TtE
Purchase price	100,000	100,000	100,000	100,000
Tax relief in year 1	0	0	20,000	0
After-tax contribution	100,000	100,000	80,000	100,000
Value of asset in year 2	108,000	110,000	110,000	110,000
After-tax withdrawal	108,000	110,000	88,000	109,000
Tax paid in year 2	2,000	0	22,000	1,000
Present value of year 1 tax relief	0	0	21,000	0
Present value of tax paid	2,000	0	1,000	1,000

In contrast, revenues with the expenditure tax come only in the second period and will depend upon actual returns. The RRA ensures government receives some revenue up front and receives a share of any excess returns.

The RRA effectively provides a tax-free allowance equal in value to the normal risk-free rate multiplied by the amount invested. Operationalizing it would create some complexities, including over the choice of the normal risk-free return, increased record-keeping requirements, and the treatment of 'losses'.

As mentioned above, the normal return can generally be well approximated by the interest rate on medium-maturity government bonds. This interest rate fluctuates, and to maintain neutrality across assets and across time, one would ideally like to ensure that the risk-free rate allowed by the tax code varied closely with it. But this clearly complicates administration, and there will always be a trade-off between varying the rate too frequently and maintaining strict neutrality.

The record-keeping required with an RRA system would be somewhat more onerous than under some other systems, but no more than under a standard capital gains tax. And there is also the question of dealing with returns below the normal rate. Giving an allowance for a normal rate of return would give rise to a tax loss when the return realized in a given year is below the normal return. The RRA allowance would then be higher than the

return it is supposed to be deducted from, giving rise to 'unutilized' RRA allowances.[22]

Unutilized RRA allowances are analogous to losses that can arise under a standard income tax, but 'losses' relative to a normal rate of return will be more prevalent than the losses in absolute terms that arise in standard income and capital gains taxes: nominal returns are below a positive rate-of-return allowance more often than they are below zero. Loss offsets are a vital aspect of the way an RRA deals with risky returns, preventing asymmetric treatment of gains and losses creating an important disincentive for risky investments.[23]

Finally, it is worth noting that the labour earnings tax, expenditure tax, and RRA approaches all achieve equal treatment of capital gains and cash income, and do not require indexation for inflation. Hence they avoid distortions to the form and timing of saving. This is immediately obvious for the EET regime. I pay tax on the value of my savings at withdrawal. It makes no difference whether they have grown as a result of accumulated interest or capital gains. The same is true for an RRA. An allowance of, say, 5% of the initial investment is carried forward. If either interest or capital gains are realized in the next period, any tax liability is set against the allowance. If capital gains are not realized until a future period, then the unused allowance is carried forward, uprated at the normal rate of return. The result is that normal returns are not taxed, whether they arise from interest or capital gains. Above-normal returns are taxed and the net present value of tax paid is unaffected by the form or timing of the returns.

13.2.3. Tax-Smoothing and Different Marginal Rates

In laying out the details of the various savings-neutral tax systems, we have so far simplified our discussion significantly by assuming that underlying tax rates are constant—a flat tax system—whereas in actual fact all modern tax systems have tax schedules with a marginal rate that is not constant.

[22] We adopt the standard terminology of 'unutilized RRA allowances' here, but it should be recognized that the 'unutilized RRA allowance' would, in fact, be more than the full RRA allowance if there were nominal losses.

[23] See e.g. Cullen and Gordon (2007).

Table 13.4. The impact of progressive taxation (40% when saving, 20% on withdrawal)

	With normal 5% return		With 10% return	
	EET	TtE	EET	TtE
Purchase price	100,000	100,000	100,000	100,000
Tax relief in year 1	40,000	0	40,000	0
After-tax contribution	60,000	100,000	60,000	100,000
Value of asset in year 2	105,000	105,000	110,000	110,000
After-tax withdrawal	84,000	105,000	88,000	109,000
Tax paid in year 2	21,000	0	22,000	1,000
Present value of year 1 tax relief	42,000	0	42,000	0
Present value of tax paid	–21,000	0	–20,000	1,000

Consider first a system in which tax rates are higher when incomes are higher—one like the UK's, with a basic rate and one or more higher rates. Suppose, in the example in Tables 13.2 and 13.3 above, that the saver is a higher-rate taxpayer in year 1 and a basic-rate taxpayer in year 2. Then the calculations for the EET and TtE systems look quite different, as Table 13.4 shows.

The EET system subsidizes saving in a way that encourages people to save at times when their tax rate is high and to access the returns when their tax rate is low. Conversely, saving would be discouraged at times when the individual temporarily faces a low tax rate. This creates a non-neutrality when the tax system is progressive. The tax system affects the level and timing of saving.

In principle, we should not seek to impose more tax on someone whose annual income fluctuates between £20,000 and £60,000, averaging £40,000, than on someone who earns £40,000 every year. But in fact at present in the UK we do. Put another way, suppose there is a threshold for higher-rate tax of £40,000. In period 1, someone earns £80,000, saving £40,000. In period 2, he earns nothing and consumes only that £40,000. Annual consumption is never above the higher-rate threshold, any more than it would have been had earnings and consumption equalled £40,000 in each period. It is not clear why the higher rate of tax should ever be payable in this example.

The example individual in Table 13.4 pays a lower rate of tax on income in the second period because his income has dropped below the threshold for higher-rate tax payments. An individual whose lifetime income is high enough to be a higher-rate taxpayer in both periods would pay more tax than the individual in our example.

Under the TtE regime, by contrast, tax at the higher rate is paid in the first period, and the fact that the tax rate is lower in the second period is immaterial (at least if only the normal return is earned). The fact that income drops does not impact on total tax paid on the savings.

This illustration shows that with non-linear tax systems, the savings neutrality underpinning the three alternative tax treatments we have considered thus far is not guaranteed. Neutrality only strictly holds if the marginal rate of tax on expenditure is constant.

In our example, saving is subsidized by a consumption tax. Equally, if the marginal rate rises as expenditure rises, then this can result in an implicit tax on saving during those parts of the life cycle in which consumption is low. Consider someone deciding whether to save some income now and spend later. Perhaps she is thinking that she might be supporting young children in the future and hence expects higher consumption needs. With constant income, this would imply that her consumption will be higher tomorrow than it is today. If there is a progressive pure (EET) expenditure tax, then she may face a higher marginal rate on her (higher) consumption tomorrow than she does on her (lower) consumption today. Consumption tomorrow is more costly—exactly the impact of a tax on savings and clearly a removal of the neutrality condition that left the timing of consumption undistorted.

Again the reverse is true with a TtE (or TEE) regime. In this case, saving in the period of lower consumption and withdrawing savings at periods of higher consumption reduces the overall tax payment.

If there is a straight choice between implementing a pure EET system and a pure TtE (or TEE) system, then it seems that, in the face of a tax system with more than one marginal rate, one group of savers must be advantaged and another group disadvantaged. If we choose EET, then those, for example, saving from high incomes now who will be facing a lower tax rate in retirement will be provided with a major incentive to save. But those wishing to save from lower income now to finance a period of higher consumption when marginal rates are higher will be disadvantaged. On the

other hand, a TtE (TEE) regime would disadvantage someone facing a higher rate now and saving for retirement when the marginal rate is lower—their average rate would be 'too high' from a lifetime perspective.

One strategy for resolving this tension would be to offer taxpayers the choice between EET and TtE regimes. This would in principle allow full 'tax-smoothing', an idea with a distinguished history in economics.[24] In the example above, if the individual with increasing expenditure could also invest in a TEE-taxed asset, this would allow her to 'smooth out' the tax rate across today's and tomorrow's consumption. Tax-smoothing permits changes in the marginal tax rate to be evened out over the lifetime.

If people have variable earnings and variable spending needs over their lives, as they surely do, then a 'pure' earnings tax (TEE) penalizes people with variable earnings. On the other hand, a 'pure' expenditure tax (EET) penalizes people with variable spending. If people had perfect foresight, faced no uncertainty in their earnings, and had complete flexibility to save (and borrow) as much as they liked within each regime, they could smooth their tax base across years, meaning that people with the same lifetime earnings would pay the same lifetime taxes (in present-value terms). This would give, in effect, a lifetime tax base—which, as we noted at the beginning of this chapter, has significant attractions.

But there are disadvantages to allowing complete flexibility. Not everyone has the foresight, the certainty, the financial sophistication, or the access to credit to smooth their tax payments perfectly. Is it fair and efficient to favour those who do? In addition, we might not want to be neutral over the timing of labour supply and consumption. For example, we might want to provide stronger work incentives around the time people are considering retirement. We will revisit this question in Section 13.3.

There is no overwhelming argument in favour of allowing any particular degree of smoothing. Even more analytically difficult is the question of how to allow for the effect of means-tested benefits (and tax credits) on saving incentives and possible gains from tax-smoothing. The analytical problem is that means-tested benefits mean that those with low incomes face higher marginal tax rates than those with higher incomes. In practice, means-tested benefits are likely to have much greater effects on saving incentives than the

[24] Dating back to Meade (1978) and Bradford (1982).

variation in income tax rates.[25] This adds great complexity, and certainly leaves no simple analytic solutions. We discuss the policy issues that arise in some detail in the next chapter.

13.3. THE ECONOMIC CASE FOR A SAVINGS-NEUTRAL TAX SYSTEM

We noted at the start of this chapter that neutrality provides a useful baseline from which to judge savings tax reform, but that the economic case for the optimality of a neutral tax system remains an active topic of research. A savings-neutral tax code has the potential to provide a tax base that raises revenue without distorting the timing of consumption. But should the timing of consumption be sacrosanct, or could some deviation from the savings-neutral benchmark be part of the optimal tax system?

As we have already seen, a progressive rate structure can undo the neutrality of a consumption tax in any case. We have also argued that our reasoning applies to life-cycle savings and not to inheritances and bequests. Moreover, we have made the case for taxing excess returns and will suggest pensions are subject to a separate tax treatment. There are several other quite sophisticated arguments for considering departures from the principle of savings neutrality across time. In considering directions for tax reform, we need to ask whether the benefits of such departures outweigh their costs, particularly given our limited knowledge of exactly what departures would be optimal.

Here we look at four arguments for deviating from neutrality.

First, it may be that people who save are more patient or have higher cognitive ability than those, with the same earnings, who don't. If patience or cognitive ability is associated with earning capacity, then taxing savings may be an indirect way of taxing those with high earning capacity.

Second, if individuals are choosing between investing in human capital and investing in financial capital, and if there are market failures that make it difficult for individuals to access credit in order to invest in their own human

[25] See Blundell, Emmerson, and Wakefield (2006).

capital, then a zero rate of tax on savings may distort decisions in favour of financial investments over human capital investments.

Third, if there is uncertainty about future earning ability, then individuals may save to hedge against a bad outcome. If the world turns out well, they will have 'oversaved', and in that situation they may reduce their labour supply. Some form of tax on savings can increase efficiency in such settings.

Finally, if future consumption (saving today) is complementary to leisure today, there may be a case for taxing savings, in just the same way as there may be a case for higher rates of indirect tax on goods and services the consumption of which is complementary to leisure (see Chapter 6).

The next subsections explore these arguments in more detail. Each is, in principle, a coherent argument against savings-neutral taxation. How far these arguments can be translated into practical policymaking will likely depend on empirical evidence that is still in the early stages of development.

13.3.1. Patience, Cognitive Ability, and Self-Control

Different people have different attitudes towards the future and towards risk. They also have different abilities to process information. As a result, saving behaviour varies. Some people save more for the future than others because they are more patient. Some save more because they have a greater understanding of the options available and the consequences of saving, or not saving. Some will be willing to bear more risk in their savings portfolios.

If it is the case that those with more patience or cognitive ability do save more, and they also have higher earning capacity, then a case for taxing the normal return to savings might emerge.

The standard argument against taxing the normal return to savings rests on the assumption that taxing savings creates inefficiencies and cannot help with redistribution. But if the observed level of saving is a good proxy for earning capacity, then taxing savings might be a useful way of redistributing. At the margin, by taxing savings the government could raise revenue and redistribute from those with higher earning capacity while reducing tax rates on labour supply and effort.

In fact, there is good evidence from experimental psychology of a significant relationship between cognitive ability and patience.[26] Those with higher ability value the future more, and will therefore save more, independently of actual earnings levels.

Higher-ability (and, particularly, more numerate) individuals are also able to process information and make complex decisions more easily. Experimental evidence suggests that they are less susceptible to framing effects (such as failing to realize that '25% fat' and '75% fat-free' mean the same thing) and generally draw stronger or more precise meaning from numbers and numerical comparisons.[27] This makes it easier for them to make rational decisions over saving and may in part explain why cognitive ability has been shown to be associated with a higher likelihood of holding stocks, and of having a private pension, even when controlling for the overall level of financial wealth and earnings.[28] This evidence also suggests that people with higher cognitive ability show more self-control.

From this sort of evidence, it is also possible to discern an additional argument for making sure that we tax any above-normal returns to savings. Higher-skilled individuals appear to be less risk averse than the lower-skilled. They may therefore invest in assets with more risk but a higher expected return. Ensuring supernormal returns are taxed can then aid efficient redistribution.

Behavioural arguments[29] such as these also suggest a case for taxing savings because saving is an indicator of having high earning capacity. The converse is that those with low earning capacity may not save enough and we might then want to subsidize, compel, or otherwise encourage them to save. In the UK, in fact, much debate has centred on apparent undersaving, and new measures have been enacted[30] that will lead to automatic default into employer-sponsored private pension saving. That is, all employees will automatically save in a pension unless they actively choose not to. The framing of the decision has been changed and the expectation (based on

[26] See the detailed evidence and arguments in Banks and Diamond (2010).

[27] See the review in Bernheim and Rangel (2005).

[28] The evidence is reviewed in Banks and Diamond (2010).

[29] Bernheim, 2002; Bernheim and Rangel, 2005.

[30] HM Treasury and Department for Education and Skills, 2007.

evidence of the effectiveness of auto-enrolment[31]) is that many more people will engage in pension saving as a result.

13.3.2. Neutrality between Financial and Human Capital Investments

People do not only have the option of investing in financial assets. They can also invest in their own 'human capital'—that is, education and skills. Such investment should earn a return in the long run, just as savings earn a return. There is no obvious reason why we should wish to distort people's choices between these two forms of investment.

In some cases, a savings-neutral treatment of human capital investment occurs in a fairly natural way. Suppose I reduce my hours of work, or delay entry into work, in order to invest in learning. No tax is levied at the time the investment occurs. The return is taxed only when income from the investment is earned—exactly the EET or consumption tax treatment we have described. (Though, just as with the consumption tax, a rising marginal rate over time may create a disincentive. If I am facing a 20% tax rate now, but that rises to 40% later as a result of my investment in human capital, neutrality is lost.)

A problem occurs if I lack access to credit. In that case, I might invest less than I would like because I can't borrow to finance consumption while undertaking the education or training. Savings-neutral taxation of financial investments may then make things even worse, since alternatives to human capital investments that reap early rewards become even more attractive. This may lead me to choose an occupation with large earnings up front rather than invest in a longer-term career that involves human capital investments.

It should be obvious why this is potentially a serious distortion. Of course, it is not just related to the savings tax system—we would like to address the lack of access to credit, but the reluctance of creditors to lend to individuals without satisfactory collateral is hard to overcome. This helps explain why so much of the costs of education and human capital investments are covered

[31] See the influential study by Choi et al. (2004).

through public provision. In the UK, as elsewhere, much formal education is free at the point of use at least until the age of 18. While this does not directly address the cost of earnings forgone, it can go some way to removing underinvestment due to borrowing constraints.

Even so, the difficulty of borrowing against prospective earning capacity means that, in practice, it is difficult to ensure full neutrality between human and financial capital. The costs of education and human capital investments should also be fully deductible for tax purposes. But many of the costs of human capital investments are difficult to measure. If it is impossible to provide a tax treatment of savings that does not distort the choice between human capital investment and financial capital investment, some taxation of the normal risk-free return from financial capital investment may be desirable.[32]

13.3.3. Earnings Risk

None of us can be certain about how much we will be able to earn in the future. We may save not just to smooth our consumption over predictable life events—having children or retiring—but also to protect ourselves in case something goes wrong—being made redundant or getting sick, for example. This will be especially true when adequate products are not available on the insurance market.

Those of us lucky enough to maintain a high earning capacity (to avoid getting sick, becoming disabled, or being made redundant) may find ourselves with more wealth than we had planned. A natural response would be to choose to work less (or less hard) and run down this unexpected wealth. The more redistributive is direct taxation, the more attractive is this option. At the margin, taxing savings or introducing an assets test—reducing access to benefits when assets exceed a certain level—will weaken the desire to take this course of action and will therefore reduce the distortionary effects of redistributive taxation.

Beyond the treatment of savings, this argument has direct application to the way we might think about making incapacity and other benefits available at later ages. It depends, of course, on the inability of governments to fully

[32] Jacobs and Bovenberg, 2008.

measure an individual's true productive ability. Imposing an assets test[33] effectively imposes a tax on savings. It reduces the incentive to build up assets and then falsely claim disability benefits.

Although the earnings risk argument is quite compelling, the practical importance of this effect will depend on how risky earning capacity is and the observability of ability (or disability).[34]

13.3.4. Interactions between Work and Savings

Finally, it is worth noting a further argument for taxing savings, which harks back to the relatively simple idea that if specific expenditures are directly related to labour supply, then they should be taxed differently from other types of consumption goods. We considered this in the discussion of the pros and cons of a uniform VAT in Chapter 6. In principle, consumption of goods or services that are complements to leisure should be taxed more heavily so as to increase work incentives.

But a similar argument will hold when we consider consumption today and consumption tomorrow. It may be efficient to discriminate against (or, indeed, in favour of) saving if, given the level of earnings, the way in which people want to divide their expenditure between consumption today and consumption tomorrow depends upon how many hours they work. By acting as a tax on future consumption, taxing savings may increase the incentive to work if consumption tomorrow is complementary to leisure today.

I may, for example, want to defer more consumption into my retirement if, during my working life, I have spent time investing in mastering and gaining pleasure from leisure activities, activities that I will be able to spend time and money on in the future. If this is the case, taxing savings may increase labour supply. On the other hand, if long hours of work today are associated with more consumption in the future, perhaps because I don't have time to consume much today, then the case is reversed and there is an argument for subsidizing saving. Whether this argument applies in practice with much force depends on empirical magnitudes that have yet to be well determined.

[33] See Golosov and Tsyvinski (2006).

[34] See Farhi and Werning (2007).

13.3.5. Taking These Arguments to Policy Design

All of these arguments are well founded in economic theory. They justify levying some tax on the normal return to capital, though not necessarily at full labour income tax rates, as in a comprehensive income tax.[35] However, there are several reasons to be cautious in applying these arguments immediately to policy.

First, in all cases, achieving the most efficient outcome requires calibrating tax rates to particular, difficult-to-measure, behavioural parameters: how much greater earning capacity savers have than non-savers with the same level of actual earnings; how much more or less complementary to work consumption tomorrow is than consumption today; and so on. In some cases (such as complementarity to leisure), it is not even obvious in which direction the departure should go—should we tax or subsidize saving?—and it is possible that some of the arguments might offset each other. Being unable to estimate the parameters of interest precisely is not on its own a good argument for persisting with neutrality: taking a best guess would presumably be better than not even trying. And as empirical economic and psychological research progresses, more accurate approximations to the optimal treatment are becoming achievable. But a degree of humility is in order: we are still some way from a robust and accurate quantitative understanding of all the relevant aspects of behaviour, and we should be mindful that a rough approximation would only yield part of the efficiency improvements that the theoretical arguments suggest might be attainable.

Similarly, even where we can be confident taxation of the normal return ought in principle to be positive, the optimal policy to achieve that may not be obvious, even theoretically. For example, a tax that reduces the net rate of return to savings (like standard income taxes) has effects that compound over time, as we discussed in Section 13.2.1. The alternative of applying an additional tax (or reduced relief) on contributions or withdrawals would also imply a net tax on the normal return (as with the case of varying marginal tax rates discussed in Section 13.2.3), but without this compounding property. Which time profile is right can be difficult to judge. Savers may be more able than non-savers, but are those who save for long periods more able than those who save for short periods? It may be that the optimal tax

[35] Banks and Diamond (2010) make this case.

profile is more complicated than either of these alternatives, or it may involve different instruments entirely, such as assets tests for disability benefits or intervention in the provision of education. The general point is again that unless the tax response can be designed to target precisely the objective identified, departing from neutrality may result in little efficiency improvement. Note again that we are considering neutral taxation for lifetime savings and not for generational transfers. We are also arguing for a system that captures excess returns.

Not only must we be cautious about how much of the potential benefits of departing from neutrality can be realized in practice; we must also be mindful of the costs of doing so. Taxing the normal return to capital as it arises inevitably reintroduces the difficulties of dealing with inflation and capital gains that plague standard income taxes. And all of the arguments are reasons for deviating from our first kind of neutrality: neutrality over the timing of consumption. They provide little, if any, rationale for departing from neutrality across assets. Yet we have already seen that a major difficulty with taxing the normal return to capital is precisely that it is difficult to be consistent across all assets. Unless the tax on the normal return can be applied consistently to defined benefit pensions, housing and other durable goods, human capital, derivatives, and so on, it will come at the cost of neutrality between assets.

A case could be made that the benefits of some (even very approximate) movement towards the theoretically superior positions described in the previous four subsections justify accepting some of the problems it would reintroduce. But taking all of the counter-arguments together, we think it would be better to make neutrality the central goal of savings tax policy.

13.4. DIRECTIONS FOR REFORM

Whilst not everyone is by any means fully 'rational' in their decisions of how much to save and in what form, it is evident that tax incentives can have a large impact on behaviour. Getting the taxation of savings right matters a lot on both equity and efficiency grounds.

The argument for taxing income from savings on the grounds that 'all income should be taxed the same' does not stand up. Saving simply defers consumption, so taxing savings means taxing earnings spent tomorrow more heavily than earnings spent today. And standard income taxes not only fail to achieve neutrality in the choice over when consumption should occur, but also make it extremely difficult to achieve neutrality between different forms of savings.

Trying to make the tax system savings neutral is a constructive benchmark. Neutrality can be achieved through a cash-flow expenditure tax (EET in our notation), a labour earnings tax (TEE), or a rate-of-return allowance (TtE). These systems differ in several important respects:

- The consumption tax and rate-of-return allowance impose a tax on above-normal returns. The earnings tax does not.
- The timing of tax payments and revenues differs, with the consumption tax involving no tax payment (revenue receipt) until consumption occurs.
- With more than one income tax rate, a consumption tax penalizes those with variable consumption, while the earnings tax and the RRA penalize those with variable earnings.

We have argued that a pure earnings tax (TEE) is inappropriate for assets with returns that may exceed the normal level. Widespread use of an earnings tax for such assets could see much effort devoted to avoidance and would allow the effort of those who are skilled at, for example, stock-picking to go untaxed.

For assets where there are likely to be significant excess returns, there is a strong case for using either a rate-of-return allowance or a cash-flow expenditure tax. Either would bring excess returns into the tax base and treat capital income and capital gains in a consistent and uniform way.

There are in fact arguments for having some element of each of the three within the tax system. As we shall see in the next chapter, in purely practical terms we are already close to an EET treatment of pensions and it would be difficult to move away from that. On the other hand, an RRA treatment of other assets may be easier to introduce in the short to medium term than a full consumption tax, not least because of transition difficulties and the timing of revenue flows to the Exchequer. But there is a principled argument for allowing a choice between these treatments in any case. Giving people the

choice allows them to 'tax smooth' so as to allow the tax system to approximate a tax on lifetime income.

But we have also seen that there are arguments against allowing full smoothing and, indeed, against a fully neutral tax treatment of savings. As far as full smoothing is concerned, the main issue is one of equity. It seems likely that only the unusually well-informed, and relatively well-off, would take full advantage. In addition, allowing full smoothing would prevent us from having tax rates that vary with age, a flexibility which we argue in Chapter 4 is potentially very valuable.

The arguments against full neutrality are rather more subtle. It may be that the decision to delay consumption tells us about someone's earning capacity. Those who are more cognitively able may be more likely to save. Savings-neutral taxation may distort decisions in favour of financial saving over human capital investment if there are credit constraints or if it is hard to measure and offset the full costs of human capital investment. It may be that taxing savings will increase the labour supply of those who have saved against the possibility of losing earning capacity but who find, *ex post*, that they didn't need to save for that reason. Or it may be that future consumption is a complement to current leisure.

These are important arguments, but we maintain neutrality as a useful benchmark as it is hard to know how to fully operationalize them from a policy point of view. Particularly given that we start with a tax system that is a long way from a tax system that is savings neutral, it seems to us to make sense to move towards neutrality. But these arguments may be enough to suggest some limits on the extent to which people can access tax-neutral savings. We therefore develop in this book a set of proposals that would provide a tax system with a neutral treatment of life-cycle savings for the vast majority of taxpayers. It would also retain—indeed, increase—taxes on capital income and gains in excess of the normal rate of return for substantial asset holdings.

Importantly, the focus of all this analysis has been on lifetime savings. The arguments for not distorting the timing of consumption for the most part concern life-cycle savings and not transfers across generations. Different arguments come into play with transfers of wealth and we dedicate the whole of Chapter 15 to the issue of how to tax inheritances and other transfers.

In the next chapter, we look in much more detail at how tax treatment currently differs between assets, and some of the practical issues in savings tax reform. We will also look at some important issues that we have barely mentioned here—including the role of means-tested benefits—and others, notably the role of National Insurance contributions, that we have avoided altogether so far.