

Taxable and tax-advantaged portfolio management for UK personal investors

R.G. Thomas

E-mail: r.g.thomas@kent.ac.uk

This paper makes some observations on the interaction of United Kingdom taxation and portfolio decisions by a personal investor managing her own investments in quoted company shares. Five holding vehicles are considered: three types of tax-advantaged account are considered (ISAs, SIPPS, and spread bets); and two types of taxable account (a company controlled by the investor, and direct holdings in the investor's personal account). Some observations are made on ways in which portfolio management of a taxable account differs from management of a tax-advantaged account. Simple models are used to illustrate the difficulty of producing post-tax out-performance from active management of a taxable account. Guidelines are suggested for allocating different types of investments across the five types of accounts, and for decisions on switching between investments in a taxable account. Several quirks are noted in the CGT legislation which are useful for the active personal investor to know. Readers who want to read just results should go directly to sections 9 to 11 of the paper.

Keywords: tax, portfolio management, personal investor, loss harvesting, Monte Carlo simulation, Wilkie model, Brownian bridges

1. Introduction

This paper makes some observations on the interaction of United Kingdom (UK) taxation and portfolio decisions by a personal investor managing her own investments in quoted company shares. Most extant literature on portfolio management, whether academic or practitioner-orientated, seems to largely ignore taxation. Similarly marketing materials from UK private client investment managers generally focus on pre-tax returns. In recent years a small literature on taxable portfolio management has developed, mainly with reference to United States (US) tax regimes¹, but this literature generally gives no consideration to tax regimes outside of the US. Although my focus is on the personal investor who manages her own investments, most of the concepts I discuss also seem potentially relevant to firms which offer investment management services to such investors. The concepts can also be applied, with suitable adjustments, to other entities which pay tax on their investment returns – for example, insurance companies, other companies, and most private trusts.

I assume that the investor is UK-resident and domiciled, and a higher-rate income tax payer who can hold shares in five main ways, each of which I refer to as “accounts”:

- (1) in a tax advantaged Investment Savings Account (ISA) to which contributions may have been made since 1999; and possibly in a fund accumulated in a Personal Equity Plan (PEP), the predecessor scheme under which contributions could be made between 1987 and 1999;
- (2) in a tax-advantaged Self-invested Personal Pension (SIPP);
- (3) in a UK-resident company which the investor controls (“corporate account”);
- (4) directly registered in the investor's own name or a nominee (“personal account”);
- (5) for short-term holdings, as spread bets.

¹ Jeffrey & Arnott (1993), Arnott et al (2001), and Berkin et al (2003) are examples of papers written from the perspective of professional investment managers. In the economics literature, there are more theoretical papers: for example Constanides (1983, 1984) considers the effect of personal taxation on capital market equilibrium; Dammon et al (2004) and Garlappi & Huang (2006) consider stylised lifetime consumption/investment models with taxable and tax-deferred accounts. In the actuarial literature, Sherris (1992) uses the Black-Scholes formula to calculate reserves for deferred capital gains tax, but this paper is not concerned with portfolio management.

I characterise (1), (2) and (5) in the list above as “tax-advantaged” accounts.² I characterise (3) and (4) as “taxable” accounts.

This paper is concerned with questions such as: which type of investments should be held in which accounts? In which accounts should turnover of investments be concentrated? How should the differing tax positions of the different accounts influence portfolio management? Some of the papers by US authors refer to the first of these questions by the term *asset location*, by analogy with the more familiar terminology of *asset allocation*. But this terminology seems less apt as a description of the second and third questions, which are also important, and so I prefer not to use it.

This paper is concerned with decisions for an investor who holds an investment portfolio through the conventional accounts listed above. For investors with substantial funds, other less conventional accounts may offer some tax advantages. For example, an offshore life insurance bond can be written on an underlying portfolio of assets under the effective direction of the investor; or an authorised unit trust can be created for a single investor or family (the so-called “private unit trust”). In my view such schemes can be effective, but they tend to exploit tax legislation in a way not intended or envisaged by legislators, and as such are vulnerable to court challenges by the authorities or legislative changes. They also tend to involve relatively high charges. I have no experience of using such schemes, and for the purpose of this paper I shall ignore them.

This paper does not consider complex tax planning schemes for single transactions, for example when an entrepreneur is considering sale of a company and steps are taken to structure the transaction in a way which mitigates the chargeable gain. The accounting and legal professions appear to devote considerable effort to developing such schemes, and the taxation authorities to challenging them. But as far as I am aware, this level of effort has not previously been applied to the quantitative and strategic issues which this paper considers.

Two current schemes which offer income and capital gains tax advantages for certain types of investment are the Enterprise Investment Scheme (EIS) and Venture Capital Trusts (VCT). However, under these schemes the nature of the underlying investments is restricted to very small unlisted or AIM-listed companies, and so the schemes are of limited use for managing a broadly based portfolio. For the purposes of this paper I shall ignore them. I shall also ignore the tax effects of death – capital gains forgiveness, and inheritance tax (unless the portfolio passes to the spouse of the investor). For older investors the prospect of CGT forgiveness may sometimes be a reason for not realising assets; and shares in trading companies listed on the Alternative Investment Market may be attractive, because they become exempt from inheritance tax after a two-year holding period. But these are subjects for another paper.

² The accounts are not entirely tax-free: stamp duty is paid on share purchases, a vesting charge may apply to the SIPP at the time of first drawing benefits if the fund exceeds the lifetime limit, betting tax is paid by the bookmaker in spread betting, etc. But the accounts are tax-advantaged in the sense that there are no income or capital gains taxes on the accumulation of a fund, which is the main focus of this paper.

2. The main tax rules

This section summarises the tax position of the various types of account, which will be well known to many readers. As stated in section 1 above, this will be on the basis that the investor pays income tax at the higher rate on earned income (which may include pension income).

Tax rates and regimes do of course change from time to time; for example, the rate of capital gains tax might be increased, reduced, or the tax abolished completely. Such possibilities need to be borne in mind when considering long-term strategies predicated on the current regime. But most changes tend to be gradual rather than revolutionary; and if changes are adverse, some degree of protection is usually given to existing rights. Furthermore, I believe that many of the *concepts* in this paper would be capable of application, with suitable adjustments, under a variety of plausible tax regimes.

Personal account

If shares are directly held by the investor with no tax wrapper (an arrangement which I shall refer to as a holding in the investor's "personal account"), the tax position for the investor can be summarised as follows. Income tax is payable at 25% of the amount of share dividends received, and at 40% on interest from bank deposits or company or government bonds. As regards capital gains, there is an annual exempt amount (£9,200 in 2007/08). Net chargeable gains (that is, gains offset by current year losses and any unused losses brought forward from previous years) below the annual exempt amount are free of tax. In general, losses cannot be set against income. Net chargeable gains above the annual exempt amount are taxed at 40%. All gains and losses are recognised for tax purposes only when they are realised.

Chargeable gains for assets held for longer periods in a personal account are reduced by taper relief. In the legislation this is expressed as a reduction to the gain to which the 40% tax rate is applied, but for many purposes the effect may be more conveniently represented as a reduction in the effective rate of tax. Accelerated accrual of taper relief applies for business assets as compared with non-business assets (the policy rationale for this appears to be to tax entrepreneurial gains more lightly than passive investment gains). Gains on business assets are taxed at an effective rate of 40% for holding periods up to one year, 20% for holding periods between one and two years, and only 10% for holding periods longer than two years. Gains on non-business assets are taxed at 40% for holding periods under three years. For holding periods between three and ten years there is a decreasing scale of tax rates, which levels out at a rate of 24% for all holding periods over ten years. Table 1 shows the effective rates of tax for various holding periods for business and non-business assets.

Shares listed on any stock exchange which appears on the list of recognised exchanges published by the tax authorities (Her Majesty's Revenue and Customs, hereafter HMRC) are always non-business assets. However, the Alternative Investment Market (AIM) operated under the auspices of the London Stock Exchange is *not* a recognised stock exchange. *Trading* companies listed on AIM are business assets, but *non-trading* companies are not. The distinction between a trading and non-trading company is relatively clear in some cases: a company which manufactures products or provides services is probably a trading company, and a company whose sole activity is investing in other companies (or in land and property) is probably not. Where a company undertakes a mix of trading and investment activities, the distinction may not always be so clear, and may need careful consideration³. However, for the purposes of this paper I shall

³ HMRC Tax Bulletins 53 and 62 give some further guidance.

use the shorthand “trading companies listed on AIM” (or sometimes “AIM shares”) to refer to AIM-listed trading companies which do qualify for the accelerated business asset taper relief. Whilst many trading companies listed on AIM are small or speculative ventures and may have relatively illiquid trading of their shares, many others are larger (and sometimes less speculative) than many small fully listed companies. From the personal investor’s viewpoint, the main difference between smaller fully listed companies and trading companies listed on AIM is that the latter offer accelerated accrual of taper relief. As will be elaborated later in the paper, this is a very important factor in practical portfolio management.

Table 1: effective rates of tax after taper relief

Holding period (complete years)	Non-business assets	Business assets
	(incl. fully-listed shares)	(incl. trading companies on AIM)
<1	40%	40%
1	40%	20%
2	40%	10% (any time after 2 yrs)
3	38%	
4	36%	
5	34%	
6	32%	
7	30%	
8	28%	
9	26%	
10	24% (any time after 10 years)	

One important qualification to “trading company listed on AIM” concept relates to dual listed companies. The list of “recognised stock exchanges” on which a listing will deprive a company of non-business asset status includes most major overseas stock exchanges. So where a trading company listed on AIM also happens to be listed on a recognised overseas stock exchange, the company will not qualify for accelerated taper relief. This is unlikely to apply to a business operating mainly in the UK, but is often relevant for mining companies, and other companies with head offices or major operations overseas.

Special rules apply to transfers of assets to spouses or civil partners, whether by way of gift or sale. In effect, no gain or loss arises on such a transfer; instead the transferee spouse assumes the transferor’s original cost, and the transferor’s dates of acquisitions for taper relief purposes⁴. Where one spouse has a lower marginal tax rate, or an unutilised annual exemption or unrelieved losses, transfers between spouses before onward sale can reduce the sum of taxes for the two spouses. However these opportunities (and views on the non-fiscal aspects of such transfers) tend to be idiosyncratic to each household, and they are not considered further in this paper. There are also special rules for transfers to other “connected persons” (a broad category which includes many relatives). Transfers to connected persons are generally deemed to be made at market value, whatever the actual consideration (if any) paid, and losses are deductible only from gains

⁴ But not for the purpose of the share identification rules, under which the transfer is treated as a date of acquisition by the transferee. (This will be relevant if the transferee also made acquisitions of the same shares on other dates, and then later sells part of the holding; the shares sold are identified by a Last-in, First-Out (LIFO) principle.)

arising on other disposals to that same connected person. Thus there may be some opportunities (and some potential pitfalls) in transfers to other family members; but these opportunities (and views on their non-fiscal aspects) again tend to be idiosyncratic to each household. These rules are not considered further in this paper.

ISAs (formerly PEPs)

The investor can also hold investments in an Individual Savings Account (ISA). Contributions of up to £7,000 (£7,200 from 6 April 2008) may be made each year. The ISA can be invested in a range of collective vehicles, and on shares listed on any recognised stock exchange. AIM shares are not permitted (unless they have a dual listing on an overseas recognised stock exchange). The predecessor scheme of Personal Equity Plans (PEPs) operated between 1987 and 1999 with broadly similar parameters. The contribution levels varied over the years. The maximum which an investor could have contributed to PEPs between 1987 and 1999 was £88,200. There were originally some minor differences in the rules for PEP and ISA schemes, but all PEPs were statutorily redesignated as ISAs with effect from 6 April 2008.

No income tax or capital gains tax is payable on investments held in an ISA, which therefore provides tax-free roll-up and the ability to vary a portfolio without tax consequences. Withdrawals can be made at any time with no tax penalty, apart from loss of the shelter from future taxation on the funds withdrawn.

Self-Invested Pension Plan (SIPP)

The investor can also hold investments in a tax-advantaged registered pension plan. Broadly speaking, a registered pension plan offers income tax relief on contributions; no income or capital gains tax is payable on the investments held in the pension account; but income tax is payable on the pension drawn in retirement (except for tax-free retirement cash of one-quarter of the total fund utilised at vesting or partial vesting). From April 2006, former restrictions on membership of multiple pension plans were relaxed to allow “full concurrency”. This means that an investor who wishes to manage her own investments in shares and other assets can establish a self-invested pension plan (commonly known as a SIPP), even if she is also a member of an employer’s pension plan.

There is also “lifetime limit” (expected to be £1.8m in 2010/11, increasing thereafter in line with RPI) on the amounts which can be accumulated at date of vesting (ie date of first drawing pension benefits) in all the investor’s pension plans. If the investor is a member of a defined benefit plan, those benefits are valued at vesting date for the purposes of the lifetime limit in accordance with a prescribed basis. Any excess of accumulated funds above the lifetime limit attract a 25% tax charge at the vesting date; the residual excess and can then be drawn as either a lump sum or as income, in both cases subject to 40% tax.. In this paper I generally ignore the lifetime limit, that is I assume that accumulated pension benefits at vesting are worth less and so the limit does not bite.⁵

⁵ This may be a less important limitation than it first appears, because the SIPP is substantially tax-advantaged even for funds which exceed the limit at the vesting date. You get tax relief on the contributions, tax-free roll-up, and an effective 55% tax ($0.75 \times 0.6 = 0.45$) at vesting. Now $0.6 \times (1.02)^{15} \times 0.45 \approx 1.0$, so for a term of 15 years an excess return of around 2%pa during roll-up is required in the SIPP to match the proceeds of taxed personal account investment. In my view 2%pa may be a low estimate of the advantage a skilful active investor can reasonably expect to derive from total absence of taxes during roll-up.

There are many further pensions rules concerning the level of tax-relieved contributions, and the timing and form of the benefits. While these further details are important in practice, they do not directly impinge on the issues which this paper considers, and so the details have been relegated to Appendix A.

Corporate account

The investor can hold investments in a company which she controls. People who are active investors may often also own a business which operates as a corporate entity. In some cases this might be a substantial enterprise; in other cases it might be a small company with a single employee (the investor) which provides consultancy services. In these circumstances it may be more tax-efficient for the company to invest at least part of any surplus profits, rather than distributing all profits immediately as taxable remuneration or dividends to the owner (whose subsequent investment of the funds will be subject to higher rate tax). Alternatively, a company can be established and funded by the investor solely for the purpose of making investments. (Of course, many investors will not own a company, in which case all references in this paper to a corporate account can be ignored; but including them allows this paper to have wider applications.)

A UK-resident company pays no tax on dividends received from other UK companies. Capital gains are computed under the CGT rules which applied to direct holdings of a personal investor before 6 April 1998. This means that instead of taper relief, gains are reduced by indexing base costs in line with the Retail Prices Index. This indexation is available only to reduce gains, not to increase losses. Losses are offset against indexed gains to give net chargeable gains, with any excess of losses carried forward to future years; in general losses cannot be offset against other trading income of the company. Corporation tax is applied to the net chargeable gains, at the rate of tax applicable to the company. For a company whose main activity is holding investments, and which has a small number of shareholders – a so-called close investment holding company – the rate of corporation tax will normally be 30% (in 2007/08). For a company whose main activity is some other trade and not investment, the marginal rates are 20% for the first £300,000 of chargeable profits each year, 33½% for the next £1.2m of profits, and 30% for profits above £1.5m.⁶ If the company funds its investments partly by borrowing, tax relief is generally available on interest paid.

3. Some observations on tax-aware portfolio management

Clarifying approaches and objectives

When the investor wishes to draw on her investments to fund current spending on consumption, the four types of account listed above offer differing ease of access. Investments held in personal account or an ISA can be drawn on at any time without penalty (except the loss of future tax shelter when funds are withdrawn from an ISA). But for investments held in a corporate account, the payment of funds to the investor may incur tax; and for investments held in a SIPP, the funds generally cannot be drawn on until age 55 (except in circumstances of ill-health). This paper ignores these differences; in other words, it is implicitly assumed that the investor's objective is to maximise the sum of wealth across all accounts, with no preference for wealth in any one account. This seems appropriate for the assumed type of investor (a higher-rate taxpayer), for

⁶ This account is slightly simplified, in that the marginal rates quoted ignore the partial restriction on smaller companies relief which arises when franked investment income (ie dividend income from other UK companies) is received; but it will suffice for our present purposes.

whom investments are likely to be in the nature of surplus funds, and any spending from investments is essentially discretionary. It would not be appropriate where the investor is saving to meet a known liability at a known future date, or where the investor has a strong precautionary motive to maximise immediately accessible funds (in which case, wealth held in a pension account might have lower utility than other wealth). However, matters such as known future liabilities and precautionary preferences seem idiosyncratic to each investor, and so less amenable to analysis in a paper like this.

Even after the above simplification, optimal portfolio management for a taxable investor is a very complex problem. Capital gains tax is triggered by the decision to sell, and so depends on the share price at that time; but it also depends on the prices at which a share was originally purchased. Furthermore because of the Last-in-first-out (LIFO) identification of share purchases when calculating taper relief in personal accounts, a record of average cost for a share is not sufficient: the full history of dates, prices and quantities purchased needs to be taken into account. If a share is sold today, a particular set of alternative investment opportunities is available; but if the sale is deferred, a different set of alternatives will be available, which may be better or worse than those available today. For an individual, selling today implies a particular tax rate (based on accrued taper relief); but if the sale is deferred, the tax rate may well be substantially different (after accrual of further taper relief). The tax position for a particular share cannot be considered in isolation, since losses realised on some shares can be used to offset gains realised on others.

It seems clear that for a taxable investor, results net of taxation are what matters; and hence that both risk and return need to be evaluated, and decisions made, on a post-tax basis. As we will see, this can be difficult, and some popular investment aphorisms suggest that it is too difficult to be worth attempting: for example, it is often said that “the tax tail should not wag the investment dog”. My own view is that whilst it is conceivable that a portfolio managed without regard for tax may also produce a good post-tax return, this approach seems most unlikely to be optimal. The intuition behind this is that tax-aware behaviour often implies an advantage which is relatively certain at the time the decision is made; on the other hand, portfolio reallocations which incur tax often crystallise a *certain* tax cost to capture a prospective advantage, which is highly *uncertain* at the time the decision is made. (For example, suppose we hold a share standing at a loss, and we are now neutral about the share’s prospects: a tax-aware decision to realise the loss and set it against other gains will generate a certain benefit. On the other hand, suppose we hold a share standing at a gain, but we now consider that a different share has better prospects: a decision to switch will crystallise a certain tax cost, for highly uncertain benefits.)

There is one advanced mathematical approach which should be briefly mentioned. In principle, optimal post-tax portfolio management can be characterised as a problem of dynamic stochastic optimisation. Given a model for share price changes and dividends, the problem could in principle be solved by evaluating all possible future paths along a decision tree. But this is very difficult, because the size of the decision tree grows exponentially with number of time periods, and quickly becomes unmanageable as soon as one considers any realistic number of shares and time periods. For personal investors, taper relief is a substantial further complication. A couple of recent papers in the operational research literature have taken this approach (Osorio et al 2004a, 2004b; DeMigeul 2005) but due to the size of the decision tree, they can consider only greatly simplified problems involving small numbers of shares and time periods. I have no experience with dynamic stochastic optimisation, but my impression from these papers is that with currently available methods, it is not yet much help: a great deal of effort and quite restrictive assumptions are required to produce rather limited results. I do not consider this approach further in this paper, but it may be a worthwhile area for future research.

Tax-advantaged accounts offer very large benefits

In my view tax-advantaged accounts such as ISAs and SIPPS offer very large benefits to the long-term investor, which are often under-appreciated. This is partly because an advantage which is small initially compounds over the years, and the large effect of compounding a small advantage is counter-intuitive to many investors. A more subtle point is that the portfolio in tax-advantaged accounts can be varied as much as one likes with no tax consequences, so that any skill the investor has can be fully exploited. The *combination* of these two effects can lead to a much higher compound return in the tax-advantaged accounts than outside them. To the investor with a long time horizon, £1 within a tax-advantaged account is worth substantially more than £1 in a taxable account. The time and effort applied to each type of account should reflect this.

A higher downside risk tolerance may be appropriate in taxable accounts

The taxation authority's role in relation to capital gains tax is sometimes characterised as providing an interest-free loan to the investor. This reflects the fact that an investor who defers realising a gain also defers payment of tax. However, the reference to an interest-free loan is an incomplete mental model, because the time value of money is not the most important point. A better model is that the taxing authority is a risk-sharing partner who puts up no initial capital, but agrees to share in any profits, and also in losses (provided the investor has other gains to offset the losses).

One consequence of this relates to the decision to “top-slice” any holding which has become too large. Because part of an appreciated asset in a taxable account is effectively owned by the government, the level at which the asset becomes too large a proportion of the investor's portfolio is higher than if the investment were held in a tax-advantaged account. This point is reinforced by the fact that the earlier a large holding standing at a profit is reduced, the earlier tax will become payable.

The presence of a risk-sharing partner is also a reason why the investor might want to allocate investments with a high downside risk to a taxable account. However, investments with a high downside risk often also have high upside risk. This presents a dilemma, because if investor is going to make any spectacular gains, it would be better to make them in tax-advantaged vehicles. The trade-off between these two ideas is discussed further in section 5 of the paper.

Diversity of tax sale costs is valuable in taxable accounts

The benefits of diversity across different assets bearing different idiosyncratic risks are widely recognised. In taxable accounts a more subtle form of diversity is beneficial: diversity of tax sale costs. The “tax sale cost” of an asset is the tax payable (or loss relief available) if the asset were sold today. Tax sale costs of different shares can vary because they have different basis costs (that is, cost as a proportion of the current value), or different accrued amounts of taper relief or indexation relief. A portfolio of shares with diverse tax sale costs provides more options for varying the portfolio at low cost – for example by taking losses early, and deferring the realisation of gains – than a portfolio with a single risky asset, or a portfolio of shares all with similar tax sale costs. This is an example of the general point that a portfolio of options (in this case, options to realise a gain or loss) is generally worth more than an option on a portfolio.

The point that diversity of tax sale costs makes a portfolio of risky assets more valuable is another reason why a taxable investor may wish to make a higher allocation to risky shares than a tax-free investor. If some of the shares produce losses and some produce gains, this may be more

valuable in a taxable account than similar returns from all of the shares. An example is given in section 4 below. (The first reason why a taxable investor may prefer a higher allocation to risky shares, which was mentioned above, is that the taxation authority acts as a risk-bearing partner to the investor in each individual investment, providing no capital but sharing in both profits and losses).

Optimal weightings for an asset depend on its holding account and tax sale costs

The optimal weighting in a particular asset depends partly on the account(s) in which it is held, and (where any of the accounts is taxable) its tax sale cost. For example, if a share held solely in tax-advantaged accounts increases in price by 200%, a larger reduction in the holding is warranted than if the same share were held entirely in a taxable account. In the taxable account, the higher the tax sale cost, the smaller the reduction in the holding should be. This is for two reasons: (i) the higher tax cost of reducing the holding (ii) in some cases, the possibility of accruing some taper relief by holding longer (this is particularly important for trading companies listed on AIM, for which substantial taper relief accrues after one and two years).

There can also be cross-effects between size of holding and tax sale costs for different shares. Suppose one share in our portfolio has been a spectacular success, growing to a very large portfolio weighting. The portfolio now has a high exposure to the specific risk of this one share, and also a high exposure to shares as an asset class. In a tax-free account, we would probably reduce the holding in this share; this would both reduce the specific exposure, and adjust the overall portfolio exposure away from shares towards cash or bonds. But in a taxable account, the tax sale cost of reducing the holding in this share will be very high. We might prefer to live with the high specific risk, and adjust the overall portfolio exposure by selling *other* shares with lower tax sale costs. Thus the presence of one greatly appreciated share (large weighting, low basis cost, high tax sale cost) can sometimes make it more attractive to sell *other* shares, compared with the position if the appreciated share was absent from the portfolio. In other words, different shares can to some extent be regarded as substitutes for one another.

4. Modelling the tax costs of turnover

Is your alpha big enough to cover its taxes?

The question heading this section was the title of probably the first US paper (Jeffrey & Arnott, 1993) to consider taxable portfolio management from a practitioner perspective. These authors were concerned with whether the gains typically generated by active management of a taxable portfolio (relative to a passively indexed portfolio) can reasonably be expected to cover the incremental tax costs of that turnover. Writing eight years later, Arnott suggested the answer: “at best, *probably not*, and if I had to give a one-word answer, it would be *no*” (Arnott, 2001). In this section I consider some simple models of corporate and personal accounts to show that these remarks are probably also appropriate for UK investors.

Modelling the tax costs of turnover: corporate accounts

To assess the costs of turnover in a taxable account, consider the following model. The investor, assumed initially to be a company, has a finite time horizon, either 10 or 20 years. The choices are (i) passive investment in the accumulation units of a unit trust index fund (ii) active management of directly held investments. The index fund charges 0.5%pa and has a low level of turnover costing say a further 0.2%pa (commissions, stamp duty on reinvestments, and bid-offer spread). Under active management, I assume charges as follows: 0.5% stamp duty on purchases,

commissions at 0.25% on both purchases and sales, 1% for bid-offer spreads, and 1% as a charge for active management. I approximate these active management costs by a total charge payable at the end of each year of 3%pa for a one-year holding period, 2%pa for a two-year holding period, and 0.67%pa for a three-year holding period. The corporate investor pays corporation tax at 30% on gains (after RPI indexation of base costs). I assume a fixed total return of 8%pa, comprised of dividends of 3%pa and share price gains of 5%pa.

Table 1 shows the excess capital return required from active management for it to produce the same proceeds as the index fund after 10 or 20 years, given holding periods of one, two or three years (that is, turnover of 100%, 50% or 33⅓% per annum). Figures are shown separately on both pre-liquidation and post-liquidation bases, where the latter allows for the tax cost of realising the passive investment after 10 or 20 years. The figures for the two time horizons are similar, with the required excess capital return slightly greater for a 20 year horizon than a 10 year horizon; this is as expected, since for a corporate investor the capital gains tax rate is independent of the holding period.

Table 1: Corporate investor: excess capital return (%pa) required for post-tax proceeds of active management with n -year holding periods to equal post-tax proceeds of unit trust index fund

1-year holding period:

	Time horizon	
	10 years	20 years
Pre-liquidation:	5.7%	5.7%
Post-liquidation:	5.1%	5.2%

2-year holding period:

	Time horizon	
	10 years	20 years
Pre-liquidation:	2.0%	2.1%
Post-liquidation:	1.5%	1.6%

3-year holding period:

	Time horizon	
	10 years	20 years
Pre-liquidation:	1.0%	1.0%
Post-liquidation:	0.4%	0.5%

Assumptions:

- Corporate investor: tax at 30% on gains after RPI indexation (assumed 2.5%pa), NIL on dividends (or notional dividends from accumulation units)
- Unit trust index fund charges 0.5%pa, with 0.2%pa cost of turnover & other expenses
- Active management charges of 1% amc, 0.5% stamp duty on purchases, 0.25% commission on both purchases and sales, 1% for bid-offer spreads)
- Total investment return 8%pa, comprised of 3%pa dividends and 5%pa capital gains
- Corporation tax on capital gains is paid immediately on realisations; for three year holding period, a further realisation is made at the end of the horizon (ie at 10 years or 20 years).

Modelling the tax costs of turnover: personal accounts

For a non-corporate account, the comparison is more complex because of taper relief. Passive investment in the index fund implies capital gains tax at an effective rate of 24% on realisation at any duration after 10 years. For active management with 100% turnover, investment in fully listed companies implies capital gains tax at a rate of 40%. This leads to required excess capital returns in excess of 5%pa, which suggests that active management is very unlikely to be worthwhile. But this may be too pessimistic, because in practice the active investor can focus on investing in trading companies listed on AIM, which implies capital gains tax at only 20% after a one year holding period, and 10% after a two year holding period.⁷ Furthermore, the active investor can also use his annual CGT allowance for the first part of each year's gains. This can substantially reduce the effective tax rate, even for quite large funds. For example, if realised gains less losses less taper relief ("net chargeable gains") are 5% of the fund in a particular year, the effective rate of CGT paid (assuming no taper relief) on a fund of £250,000 will be only about 11% of net chargeable gains; on a fund of £500,000, about 25%; and on a fund of £1m, about 33%. Table 2 shows the range of effective CGT rates for four sizes of funds and three levels of net chargeable gains in any year.

Table 2: effective rates tax on net chargeable gains, at different wealth levels

Net chargeable gains * (as percentage of end of year pre-tax fund)	Total fund at year end (outside of tax-advantaged wrappers)			
	£100,000	£250,000	£500,000	£1,000,000
5%	NIL	11%	25%	33%
10%	3%	25%	33%	36%
15%	15%	30%	35%	38%

(*) Net chargeable gains = realised gains less losses less taper relief
Assumptions: annual exempt amount = £9,200 (as in 2007/08).

Table 2 suggests that a wide range of effective tax rates may be relevant to the personal investor, particularly at lower levels of wealth held outside tax-advantaged accounts. We also need to remember that the annual exemption cannot be carried forward from year to year, and that the current year's losses must be set against gains before applying the exemption; this means that in some years, the exemption will be partly or wholly wasted.⁸ Taking all this into account, an effective CGT rate of 30% seems a reasonable illustrative figure to work with, for a fund comprising investments with no taper relief.

In practice the active investor can achieve lower effective rates of CGT than this, by focusing the portfolio on trading companies listed on AIM. These qualify as business assets, and so accrue 50% taper relief one year after purchase, and 75% after two years, which implies effective CGT rates of 20% and 10% respectively. However, before assuming these lower rates, we also need to note that the benefit of taper relief is often reduced by the order in which loss relief, taper relief and the annual exemption are applied. This is determined as follows:

⁷ Focusing mainly on AIM companies also has the very important advantage that these smaller companies tend to be less well researched than larger companies, and so offer more opportunity to the active investor.

⁸ The position in respect of earlier years' losses is slightly more favourable: these can be carried forward to set against gains in a later year; and the annual exemption in each later year can be used in priority to the brought forward losses, which can then be carried forward to the next year, and so on.

- (1) first, deduct losses from gains with the lowest percentage entitlement to taper relief, relieving each gain to a minimum of zero (any unutilised current year losses are carried forward for use in future years);
- (2) if the sum of relieved gains from (1) exceeds zero, deduct any losses brought forward from earlier years from the non-zero relieved gains with least entitlement to taper relief, but reducing the sum of relieved gains only to the level of the annual exemption (any losses still unutilized are carried forward again);
- (3) deduct any applicable taper relief from each relieved gain;
- (4) finally, deduct the annual exemption from the sum of (3), to a minimum of zero. Any part of the annual exemption which is unutilised is lost: it cannot be carried forward for use in future years.

This sequence means that taper relief is slightly less advantageous than it first appears: in effect losses as well as gains are tapered, and the priority of loss relief together with the inability to carry forward an unutilised annual exemption often mean that some taper relief is “wasted.” So even if we assume that all investments are trading companies listed on AIM with a holding period of one year, it would be too optimistic to halve the 30% average effective tax rate which we suggested above for investments with no taper relief. Taking this into account, an average effective tax rate of 20% seems a reasonable illustrative figure, for an investor who focuses his investments outside tax-advantaged accounts predominantly on trading companies listed on AIM.

I therefore show figures for a one-year holding period, with an assumed average effective tax rate of 30% as representative of a fund of fully-listed shares, and 20% as representative of a fund with a primary focus on trading companies listed on AIM. For a two-year holding period, the investor in trading companies listed on AIM will sometimes be able to take advantage of a tax rate as low as 10%; we approximate this by an average effective tax rate of 15%, and show figures for this as well as for 30%.

Table 3 shows the excess capital return required from active management for it to produce the same proceeds as the index fund after 10 years, given holding periods of one, two or three years (that is, turnover of 100%, 50% or 33⅓% per annum). The required excess capital returns are shown on both pre-liquidation and post-liquidation bases, where the latter allows for the tax cost of realising the passive investment after 10 years. A higher proportion of total return from gains rather than dividends reduces the required excess capital return for active management to be out-perform. Figures for 20 years are not shown, because at the modest rates of capital growth assumed here, any deferral of realisation after ten years (when maximum taper relief has accrued) makes very little difference.⁹

Table 3: Personal investor: excess capital return (%pa) required for post-tax proceeds of active management to equal post-tax proceeds of unit trust index fund

1-year holding period (time horizon 10 years):

	Effective CGT rate	
	30% (fully-listed shares)	20% (focus on AIM shares)
Pre-liquidation:	5.7%	4.4%
Post-liquidation:	4.4%	3.3%

⁹ For higher rates of return, the deferral of realisation becomes slightly more beneficial; but even for 20%pa returns, the required excess capital return for a 20 year horizon is only about 0.3% higher than for a 10 year horizon.

2-year holding period (time horizon 10 years):

	Effective CGT rate	
	30% (fully-listed shares)	15% (focus on AIM shares)
Pre-liquidation:	2.7%	2.0%
Post-liquidation:	2.1%	1.6%

3-year holding period (time horizon 10 years):

	Effective CGT rate	
	30% (fully-listed shares)	15% (focus on AIM shares)
Pre-liquidation:	1.8%	1.0%
Post-liquidation:	0.9%	0.1%

Assumptions:

- Personal investor pays tax at 25% on dividends, and effective 30%, 20%, or 15% (as shown) on gains (reductions from 40% approximate the benefit of annual exemption, and faster taper relief on trading companies listed on AIM)
- Unit trust index fund charges 0.5%pa, with 0.2%pa cost of turnover & other expenses
- Active management charges of 1% amc, 0.5% stamp duty on purchases, 0.25% commission on both purchases and sales, 1% for bid-offer spreads
- Total investment return 8%pa, comprised of 3%pa dividends and 5%pa capital gains
- Dividends received at end of each year, and reinvested after tax
- Capital gains tax is paid immediately on realisations; for three year holding period, a final realisation is made at the end of 10 years.

A sensitivity test for lower dividend yields

Both Tables 1 and 3 assume a typical split of returns between dividends (3%pa) and capital gains (5%pa). If a greater proportion of the total returns were received in the form of capital gains, the tax costs of turnover would be higher, and so the required excess capital return for active management to keep up with passive management would be higher. To assess the impact of this, I reworked the tables with the same total return of 8%pa, but comprised of dividends of 1%pa and capital gains of 7%pa. The increase in required excess capital return was almost always less than 1%pa, and usually less than 0.5%pa. As this is small, I have shown figures only for the typical split of returns.

Observations on Tables 1 and 3

The striking feature of Tables 1 and 3 is that in many cases the required excess capital returns are well in excess of reasonable expectations of active management. Few active managers would claim expected excess capital return over 3%pa; moreover, *any such claim is uncertain, but the tax and transaction costs of attempting to generate excess capital return are certain*. Even for lower turnover, say less than 50%pa (that is an average holding period longer than two years), the investor needs to be confident he can generate excess returns of 1-2% to keep up with a passive investment.

There is one situation where a stronger case might be made in favour of active investment. For the hobbyist investor who manages her own portfolio without paid advisers, and regards the time

she spends on this as free, our assumed 1% annual management charge on active management can be ignored. The required excess capital return with a two year holding period and a focus on trading companies on AIM then falls to around 1%. This may be a reasonable expectation for the dedicated investor allocating a modest total fund. Nevertheless, the overall impression from the figures is that generating post-tax out-performance by active management of a taxable portfolio is quite difficult.

How should the investor respond to this difficulty? For many investors an index fund, or some other collective investment with very low charges, may be a good solution. Alternatively, if the investor still wishes to pursue active management, turnover should be concentrated in tax-advantaged accounts such as an ISA or SIPP. Turnover in taxable accounts should as far as possible be restricted to transactions with unusually benign tax consequences, such as “loss harvesting” (this is explained later in the paper).

The results in this section are broadly consistent with studies based on US tax rates (Jeffrey & Arnott, 1993; Arnott, Berkin & Ye, 2001a) which suggest that an excess return of 2-3% is needed in a conventionally managed active portfolio to match the after-tax returns of an index fund.

5. Which type of assets should be held in taxable and tax-advantaged accounts?

Shares or bonds in tax-advantaged accounts?

There is an academic literature which suggests that a corporate sponsor of a defined benefit pension fund should prefer bonds rather than shares to be held by the pension fund, because of a tax arbitrage effect (Black 1980, Tepper 1981, Ralfe et al 2004). Holding bonds in the pension fund (on which untaxed interest is received) enhances the capacity of the company to issue its own bonds (on which the interest paid is tax-deductible). However this prescription is widely ignored in practice, and not universally accepted in theory. For example, some authors say that it ignores the value of the firm’s option to default on the pension obligation, which should encourage the sponsor to direct that some shares be held in the pension fund (Bicksler & Chen, 1985; Blake & Khorasaneh 2004).

For the personal investor, a similar argument in favour of allocating bonds (or perhaps shares which pay high dividends) to tax-advantaged accounts can be made. Interest receipts from bonds and dividends from shares held in the investor’s personal account are taxable, at 40% and 25% respectively, with no allowances¹⁰. On the other hand, capital gains can be wholly or partly offset by the annual exemption (£9,200 in 2007/08) for capital gains. This is an argument for holding corporate bonds or high-yielding shares in tax-advantaged accounts, and other shares in taxable accounts.

A further argument can be made relating to the optional nature of capital gains tax. The tax is optional in the sense that it crystallises only on sale of a particular asset, when either a gain or an allowable loss is created. Allocating assets with more dispersed returns (such as risky shares) to the taxable account, and assets with more steady returns (such as bonds, or high-yield shares) to the tax-deferred account may be advantageous because it facilitates use of the annual exempt amount, and also “loss harvesting.” Investments which have increased sharply in value and become overvalued (“winners”) can be sold. Tax on these sales can be minimised or eliminated by selling other investments which have fallen (“losers”), that is “loss harvesting”. In this way net chargeable gains can be reduced towards the level of the annual exemption (one should try to

¹⁰ Recall that we assumed at the start of the paper that the investor is a higher rate income tax payer.

avoid reducing them below this, since the exemption is then “wasted”). This suggests that the value of the option might be increased by holding investments with greater dispersion of returns – for example, risky shares – in the taxable account.

This intuition can be confirmed by a simple three-asset, two-period model as shown in Table 4. The overall portfolio return is an identical 20% in each period, for both portfolios A and B. Two out of the three shares in each portfolio are sold and reinvested at the end of the first period, and the whole portfolio is sold at the end of the second period. In portfolio A (with low dispersion of returns) the investor buys three shares which each rise 20% in period 1. Two of the shares are then sold, and tax is paid at 40% on the resulting gains. The third share is held. In portfolio B (with high dispersion of returns), the overall portfolio return in period 1 is 20% is comprised of very dispersed returns on equal initial investments in the three shares: 80%, -50%, and 30%. At the end of period 1 the first and second shares are sold (profits are taken, and losses harvested), and tax is paid at 40% on the resulting net chargeable gain. The third share is held.

In period 2, all shares return 20%. At the end of period 2 all shares in both portfolios are sold, and tax paid on gains. The terminal post-tax portfolio value is 3.68 in portfolio A (less dispersed returns), but 3.78 in portfolio B (more dispersed returns).

Table 4: the benefits of dispersed returns in a taxable account

	Portfolio A (low dispersion)			Portfolio B (high dispersion)		
	1	1	1	1	1	1
Investments at time 0						
Portfolio return in period		20%			20%	
Return on individual shares in period	20%	20%	20%	80%	-50%	30%
Pre-tax time 1 value	1.2	1.2	1.2	1.8	0.5	1.3
Action at time 1	Sell	Sell	Hold	Sell	Sell	Hold
Net realised gains on portfolio		0.4			0.3	
Tax payable at time 1		(0.16)			(0.12)	
Investments in new (or existing) shares at start of period	1.12	1.12	1.2	1.09	1.09	1.3
Growth in period	20%	20%	20%	20%	20%	20%
Pre-tax time 2 value	1.344	1.344	1.44	1.308	1.308	1.56
Action at time 2		Sell all holdings			Sell all holdings	
Tax payable at time 2	(0.0896)	(0.0896)	(0.176)	(0.0872)	(0.0872)	(0.224)
Net terminal proceeds		3.68			3.78	

Note that if the low-dispersion portfolio A and the high-dispersion portfolio B had each been “bought and held”, with no sales at time 1, the net terminal proceeds at time 2 would have been 3.79 and 3.47 respectively. 3.79 is very slightly better than either of the results with intermediate sales; this illustrates the point that if a share is going to continue rising, it must be optimal to defer crystallising any tax. 3.47 is considerably worse than either of the results with intermediate sales; this illustrates that for value to be realised from the optionality created by dispersed returns, some intermediate sales need to be made.

Note that the point made here relates to *dispersion* of returns from different shares in the portfolio, *not* the volatility of the overall portfolio return. If the portfolio were comprised of a single share each year which is sold at the end of the year, greater volatility would provide no tax advantage; and the greater volatility would also act as a drag on pre-tax expected logarithmic return, and hence long term compounding. It is *dispersion* of returns on different assets which is potentially advantageous in a taxable account, not volatility of the overall portfolio return.

Based on arguments along the lines above, some papers by authors considering US tax regimes (eg Shoven & Sialm, 2003; Dammon et al, 2004) derive results suggesting that the investor should allocate assets with low volatility and high yields taxable at high rates – typically bonds – mainly to tax-advantaged accounts; and allocate assets with high volatility and low yields taxable at low rates – typically shares – mainly to taxable accounts. However, as with the Black/Tepper “pension fund in bonds” strategy for companies, these prescriptions are widely ignored in practice, and disputed in theory. Under US tax regimes, tax-advantaged accounts are generally pension plans, and so the investor has only limited access to the funds before retirement age. Some authors have suggested that allocating both shares and bonds to taxable accounts may then be justified by a precautionary motive, that is a desire to smooth returns in the taxable account which can be drawn on without restrictions before retirement age (Amronin, 2003; Garlappi & Huang, 2006). In the UK, this is less convincing, because some types of tax-advantaged accounts (ISAs) can be drawn on at any time. To an active UK investor, a more compelling argument for allocating shares to tax-advantaged accounts is that tax benefits will be maximised if the largest gains from the investor’s “best ideas” are located in tax-advantaged accounts. For most active investors, most of the time, “best ideas” will tend to be specific shares rather than bonds. Also, the active investor may have an active investment strategy from which she thinks she can generate pre-tax excess returns, but which requires a high turnover. Such a strategy may well produce negative excess returns on post-tax basis in a taxable account, and so can be exploited only in a tax-advantaged account.

For the UK personal investor, which effect is more important – the income tax saving by placing bonds in the tax-advantaged account, or the capital (and income) taxes saving by placing shares in the tax-advantaged account? Intuitively, for a successful investor, the saving in capital taxes seems likely to be more important. Poterba et al (2000) show that allocating equity mutual funds in priority to bonds to a tax-advantaged account would have accumulated greater wealth under US tax rules, based on actual returns in US financial markets over the period 1962-1998. Under US tax rules, the mutual fund investor is taxed on an annual basis on the fund’s returns, but under UK rules, capital gains tax for a unit trust investor is deferred until the investor sells the units; hence the result under US tax rules is not directly applicable to investment in unit trusts in the UK. For a UK-resident investor, the form of investment with a tax status which corresponds most closely to that of US investor in mutual funds is direct investment in UK shares. On this basis, my experience accords with the result of Poterba et al: the capital gains tax avoided by holding “best ideas” in tax-advantaged accounts tends to be much larger than any advantage which might conceivably be obtained from receiving income on bonds (or shares with higher yields) tax-free in those accounts, or from increasing the value of CGT options by allocating volatile assets to

taxable accounts. Furthermore, my experience is that the ability to vary the portfolio with no tax consequences allows a much higher compound return to be obtained on a “best ideas” share portfolio in a tax-advantaged account than on broadly similar portfolios held in taxable accounts.

In conclusion, I believe that for an active investor, by far the most important consideration is to allocate one’s “best ideas” to tax-advantaged accounts. I believe that this should take precedence over sheltering high yields in tax-advantaged accounts, or increasing the value of CGT options by allocating volatile assets to taxable accounts.

6. In which accounts should turnover be concentrated?

In general, turnover should be concentrated in the tax-advantaged accounts, and any turnover in taxable accounts minimised. Any share which is expected to be held only for a short time should, if possible, be allocated at the time of purchase to a tax-advantaged account.

However, there are a couple of qualifications to this. First, *some* turnover is required in the direct personal holdings to utilise the annual CGT exemption, and to practise “loss harvesting.” Second, some turnover can have tax advantages in a company where it involves selling shares which have recently paid a substantial special dividend. These concepts are discussed below.

Loss harvesting

In general terms, loss harvesting refers to the concept of realising losses in a portfolio, while deferring the realisation of gains. Consider a taxable account where the investor makes no attempt at active stock selection, choosing the initial portfolio either at random or by reference to a market index. In a loss harvesting strategy, the investor reviews the portfolio say once a month and sells any shares which are standing at a capital loss, and reinvests in shares with similar risk and return characteristics.¹¹ If the losses generated can be set against gains in the passive portfolio arising from corporate events such as takeovers or rebalancing towards the index, loss harvesting can generate an advantage over entirely passive management. Alternatively, the losses generated can be used against the investor’s gains on assets held outside the passive portfolio: for example, the investor may have large shareholdings in her own business with a very low (perhaps nil) basis cost. The tax benefit from loss harvesting is usually more certain than any excess return which a portfolio manager might hope to generate from share selection.

Monte Carlo simulations of loss harvesting of a portfolio based on the S&P 500 index have been described (Arnott, Berkin & Ye, 2001b; Berkin & Ye, 2003). The firms of these authors (and a few other firms in the US) offer “indexed funds with loss harvesting” as an investment strategy for taxable investors. As an active investor, I would not follow such a mechanical loss harvesting strategy myself; but I found the simulations interesting as a quantitative measure of the potential benefits from tax-aware portfolio management. I have therefore carried out my own simulations using parameters appropriate to UK investors. These are described in Appendix B.

The calculations in Appendix B evaluate loss harvesting as a “pure” strategy for an otherwise passive investor. But loss harvesting does not have to be practised in this mechanical way. The general concept of sales at a loss generating a tax advantage can be useful in some contexts to a more active investor, for example as described in the next paragraph.

¹¹ In practice a loss is harvested only if the loss is sufficient to cover round-trip transaction costs on the sale and reinvestment into new shares.

Sales after special dividends

Some turnover can be advantageous in a corporate account where it involves selling shares which have recently paid a substantial special dividend. Suppose our own UK-resident company (“the investor company”) holds shares in another company (“the investee company”) priced at 100p, on which a special dividend of 50p is declared (typically because the company is returning funds to shareholders following sale of a major part of the business).¹² Provided that the investee company is UK-resident, the special dividend is franked investment income, and therefore no tax is paid by the investor company on its receipt. At the opening of the market on the day the share goes ex-div in relation to the special dividend, the share price of the investee company will typically fall by around 50p. Selling the share after the ex-div date gives rise to a substantial capital loss, which can be used against other gains. If the rate of corporation tax applicable to the company is 30%, there is an effective gain of 15p.

Some commentators (eg Arnott et al, 2001) suggest that a taxable investor should seek to build a portfolio with which he can be happy on very low turnover. The above example illustrates that this “minimise turnover” objective is not always correct: turnover in certain situations can have beneficial tax effects. The investor’s objective, correctly stated, is not to minimise turnover, nor to minimise taxes, but rather to *maximise post-tax returns*. However, I think it is probably true that maintaining some bias towards inactivity and low turnover is helpful, for two reasons. The first reason is that doing anything always implies other transaction costs apart from tax. The second reason is behavioural and epistemological: most active investors tend to be *too* active, and overconfident in their opinions about future relative returns from different assets. When you think you know something which justifies incurring transaction costs, you are likely to be wrong more often than you expect. For most investors, acting only when you are very sure indeed can lead to better results.

Loss harvesting and the fiscal calendar

Loss harvesting is only beneficial if the investor has gains which can be offset by the harvested losses. Sometimes the investor will be contemplating realising a loss in the earlier part of the tax year, without yet knowing if sufficient gains will be realised later in the year for the harvested loss to offset. In these circumstances, it is probably better to have a bias towards conservatism and inactivity. This is because the loss harvesting incurs a *certain* transaction cost (brokerage and market impact) for an as-yet *uncertain* tax benefit.

On the other hand, towards the end of the tax year, this year’s gains available to be offset by losses are relatively certain, but *next* year’s gains are uncertain. The investor may then wish to accelerate the harvesting of losses before the end of the tax year, even if he believes that the shares sold have very strong short-term prospects, as in the next example.

Example

Suppose the investor originally bought a share at 100p, but in March 2007 it has fallen to 10p. (Of course it would have been better never to buy (and hold!) such a share, but most investors occasionally make such mistakes.) The investor believes that the price fall is grossly over-done,

¹² Alternatively, the income shares of split capital investment trusts sometimes offer the prospect of a relatively predictable stream of very high dividends, followed by a less certain (but probably large) loss on redemption. These shares are often interesting to the corporate investor. Adams & Clunie (2006) provide further discussion.

and that the share is likely to recover to say 20p in the next few months. The investor also already has substantial net chargeable gains (with no taper relief) for the 2006/07 tax year.

If the investor sells for 10p now, the *certain* outcome (including the benefit of tax losses, and ignoring transaction costs) is $90 \times 0.4 + 10 = 46\text{p}$. If he waits until the new tax year and sells for 20p, the outcome will be $80 \times 0.4 + 20 = 52\text{p}$. That is, if the share price rises 100%, the outcome for the investor will be only 13% higher. Furthermore, next year the investor might not have sufficient chargeable gains to utilise the loss, so that its benefit would be deferred to a later year. Alternatively, the investor's available gains next year might happen to be subject to 50% or 75% taper, giving outcomes of $80 \times 0.2 + 20 = 36\text{p}$ and $80 \times 0.1 + 20 = 28\text{p}$ respectively. These outcomes are 22% and 39% *less* than the *certain* outcome if we sell the share immediately. Weighing up all these and similar possibilities, the investor might well decide to sell the share now, despite her belief that there is a good chance the price may soon double.

7. Developing guidelines for switching from one share to another

Capital gains taxes as transaction costs

In principle, it seems to me that any decision to vary a portfolio should be based on 'opportunity cost' comparisons, including tax effects. By this I mean that any decision to transact in a particular quantity of a particular share should be evaluated against the full set of possible transactions and tax consequences (including doing nothing, and waiting for a better opportunity) at that time. This is easier to state than to apply in practice. Taxes on income are relatively easy to deal with, because they can be netted off a future income stream in perpetuity. Taxes on capital gains cannot be treated in this way; they need to be thought of as (sometimes very large) transaction costs.

Judicious lethargy

The large tax transaction costs which are often faced by a taxable investor suggest that the level of activism in investment decisions should be lower than in a tax-advantaged account. Even if we thought we had a model which identifies an optimal portfolio on a post-tax basis, the certainty of tax and other costs of varying the portfolio and the uncertainty about the model mean that there is probably a broad band around the optimal portfolio in which the investor should take no action. In other words, judicious lethargy can be an important trait for a taxable investor, particularly when many shares in a portfolio have high tax sale costs.

An exception to the merits of judicious lethargy is where tax sale costs are negative. This applies when a share in a taxable account is standing at a loss, so there will be a tax benefit of selling (assuming that other gains are available against which to set the loss). In these circumstances, one should be more "trigger-happy" about selling the shares than if they were held in a tax-free account.

The concept of judicious lethargy leads to a curious corollary: an investor who is prepared in principle to make concentrated bets on her views may for tax reasons hold a portfolio in which he no longer has any strong views about many of the largest shareholdings. Consider an investor starting from cash with bold views which lead her to make concentrated bets on individual shares. If the concentrated bets are successful, after some years the investor's portfolio may be dominated by some large shareholdings with very low basis costs and hence high tax sale costs. The certain tax sale costs mean that the investor tends to retain these large shareholdings, even though she no longer has any strong belief that they will generate future pre-tax excess returns.

Only a *very* exceptional new investment opportunity can justify incurring the certain and high tax sale costs.

The idea of a broad region of inactivity in a taxable portfolio dependent on factors such as basis costs, duration since shares were purchased, and available CGT exemptions or losses, sits comfortably with the practicalities of managing investments in multiple accounts with differing tax positions. In particular, one very important practicality is liquidity. Investment and financial theories generally seem to pay little attention to liquidity, but to a practical investor in anything but the very largest companies it is a hugely important consideration: many shares have to be bought and sold when one can, rather than when one wants to, and it can sometimes take many months to buy or sell a holding. If the investor controls a variety of taxable and tax-advantaged accounts (for herself, for other family members, or perhaps for clients), the accounts are likely to have some common holdings. The differing tax positions often provide a rationale for which account to start with when one begins to buy or sell a particular share. For example, when buying a share which one expects to hold for less than one year, one should probably allocate it as far as possible to tax-advantaged accounts. When selling a share, sales should generally be made first from the accounts where the tax cost of selling is lowest (this will usually mean a tax-advantaged account, except where a taxable account holds shares standing at a loss).

Switching from one share to another

What limits can we place on the principle of judicious lethargy? One limit is that if the anticipated holding period were long enough, it would always be right to switch to an asset with a higher expected return, no matter how high the tax sale cost associated with the original share. This is because differences in expected return have an exponential effect on terminal wealth, but the tax payment (a transaction cost) has only a multiplicative effect on terminal wealth.

This observation can be refined for shorter time horizons. Consider first investments held by a company, for which the single tax rate and absence of taper relief makes the arithmetic a little simpler. To address shorter time horizons, it is necessary to consider tax sale costs. In the absence of taper relief, these depend only on the indexed basis cost of the old share (that is, the original purchase price, with indexation if a gain is in prospect, as a fraction of the current price). The investor also needs to have some idea of the time horizon over which she expects the new share which is bought to out-perform the old share which is sold, and also the time horizon for any ultimate liquidation of the portfolio. Note that the new “share” could be another share, or it could be a bond or cash.

Table 5 shows the excess capital return required from the new share relative to the old share to justify a sale subject to 30% capital gains tax, for various basis costs expressed as a decimal of the original purchase price. For simplicity dividends have been ignored (which can be justified either by the dividends actually being nil, or by assuming that the *difference* in the dividend streams for the new and old shares is small enough to ignore.) The figures towards the top left of the table show that where a share already held has a low basis cost, and the out-performance of the alternative asset is expected to persist for only one year, a difference of 32%pa in expected returns is required to justify a switch to an alternative asset. The figures towards the bottom right of the table show that where an existing share is held with a high cost basis and the time horizon for the future difference in expected return is long, only a small difference in expected return is required to justify a switch.

Table 5: Corporate account: excess capital return of new share relative to old share required to give equivalent fund *before* liquidation tax at end of out-performance time horizon

Indexed basis cost of old share	Out-performance time horizon (years)				
	1	2	3	4	5
0.2	32%	15%	10%	7%	6%
0.3	27%	13%	8%	6%	5%
0.4	22%	10%	7%	5%	4%
0.5	18%	8%	6%	4%	3%
0.6	14%	7%	4%	3%	3%
0.7	10%	5%	3%	2%	2%
0.8	6%	3%	2%	2%	1%
0.9	3%	2%	1%	1%	1%

Assumptions:

- Tax 30% (main corporation tax rate)
- Prospective tax on liquidation is ignored
- Excess capital returns is defined as $((1+n)/(1+o) - 1)\%$, where n is the capital return on the new share and o is the capital return on the old share.

Table 5 ignores tax after the end of the expected period of out-performance. To consider this, it is necessary for the investor to specify a terminal time horizon over which she wishes to target wealth after realization of investments, say 20 years. We can then assume that after the period of expected out-performance, the two shares compound at the same rate until the time horizon, when all shares are sold and tax on liquidation is paid. Table 6 shows the excess capital return required over various horizons from 1 to 5 years from the new share relative to the old to justify a sale subject to 30% capital gains tax, for various base costs, using a terminal time horizon of 20 years. The capital return on the original share is assumed to be 5%pa, and the same assumption is made for the identical capital return on the two shares after the period of out-performance. [Need to do a graph, a set of curves, here to illustrate the non-linear dependency on return on the original share.]

Table 6: Corporate account: excess capital return from new share relative to old share required to give equivalent fund *after* liquidation, with terminal time horizon of 20 years.

Indexed basis cost of old share	Out-performance time horizon (years)				
	1	2	3	4	5
0.2	20%	9%	6%	5%	4%
0.3	17%	8%	5%	4%	3%
0.4	14%	7%	4%	3%	3%
0.5	11%	5%	4%	3%	2%
0.6	8%	4%	3%	2%	2%
0.7	6%	3%	2%	2%	1%
0.8	4%	2%	1%	1%	1%
0.9	2%	1%	1%	0%	0%

Assumptions:

- Tax 30% (main corporation tax rate)
- Fund liquidated and tax paid at time horizon of 20 years
- Excess capital returns are measured as $((1+n)/(1+o) - 1)\%$, where n is the return on the new share and o is the return on the old share.
- Old share compounds at 5%pa during the out-performance period; after this, both shares compound at 5%pa.

The figures in Table 6 show a similar but lower pattern to Table 5. This is as expected. In Table 5, the portfolio incurred CGT if switched, but over the time horizons considered there were no realisations in the original portfolio; consequently a very large excess return on the new share was required to overcome this disadvantage. In Table 6, only the switched portfolio incurs CGT now; but CGT for both original and switched portfolios is deducted at the terminal time horizon, and that for the original portfolio is calculated using a lower base cost (because the base cost wasn't raised by an earlier realisation). This method of evaluation leads to a lower excess return required from the new share to justify the switch.

The figures in Table 6 have an interesting sensitivity to the assumption (say c) for the common capital return on the two shares after the out-performance period. For $c=0\%$, a very small excess return in the out-performance period is sufficient to justify a switch. As c increases, the excess return required to justify a switch increases; but after a certain value of c (about 10%pa for the 20-year time horizon as in Table 6), further increases in c have little effect on the required excess return. This can be understood by noting that the switch gives three effects –

- (i) a loss of part of the fund to tax paid at the time of the switch;
- (ii) an increased rate of return (but applied to a smaller start-of-period fund) during the out-performance period; and
- (iii) a reduction in tax paid (because of increased base cost) at the final time horizon on ultimate liquidation of the fund.

Then note that the fund which the switcher needs to “catch up” by the end of the out-performance period is the proceeds of investment in original share, *less* the tax saving in (iii) discounted at the rate of return c . As c increases, the discounted tax saving in (iii) decays exponentially, and so beyond a certain level further increases in c have very little effect on the “catching up” required over the out-performance period. (Do a diagram here later.)

The figures in the tables make no allowance for the point that any switch incurs a *certain* tax cost for an *uncertain* advantage in expected return. Switching also incurs other *certain* costs, including commissions, stamp duty and bid-ask spreads. This point is particularly important because most investors tend to be over-confident about their expectations and give confidence intervals which are too narrow. As already noted, a practical response to this point is judicious lethargy.

Taper relief on AIM: should one wait for one year or two before switching?

For personal account investments in trading companies listed on AIM, substantial taper relief accrues rapidly: the effective tax on gains for a higher rate taxpayer falls from 40% to 20% after one year, and then to 10% after two years. This means that the investor is often faced with a choice between switching subject to a higher rate of tax now, or waiting for a halving of the tax rate in a few weeks or months from now. In either case, payment of the tax is deferred: the tax is payable between 10 (minimum) and 22 (maximum) months later. This deferred payment means

that the funds available for investment now are not affected by the decision to switch. This facilitates a simple analysis, as follows.

Let b be the basis cost (as a fraction of the current share price) for the original share, and let t be the effective tax rate if it were sold now (so $t = 40\%$ if the original share was bought less than 1 year ago, and $t = 20\%$ if it was bought between 1 and 2 years ago). Assume that the original share can either be sold now and reinvested in a new share with higher expected return, or held until accrual of the next taper relief threshold, after which the effective tax rate will be $t/2$. If we switch, we will incur tax now; we also need some convention to account for the prospective tax liability when we come to sell the new share. The convention we adopt is to assume that whatever we are holding (either the original share, or a new share) will be sold exactly one year from now.¹³ This gives us two cases to consider:

- (i) If the original share was bought less than one year ago ($t = 40\%$), the sale 12 months hence of either the original share or the new share will be subject to tax at $t/2$
- (ii) If the original share was bought between 1 and 2 years ago ($t = 20\%$), the original share will be sold 12 months hence subject to tax at $t/2$, but the new share will be sold subject to tax at t .

(i) Original share bought less than 1 year ago (so sale in 12 months of new share taxed at $t/2$)

Let the expected capital return on the old share to the taper relief threshold be o . Let the required capital return on the new share over the same period be n . Let the basis cost of the old share be b (for example, $b = 0.5$ corresponds to a share which has risen 100% since we bought it). We ignore all dividends and discounting of tax payments.¹⁴

If we stick with the old share, the expected post-tax proceeds after one year are

$$(1+o) - ((1+o) - b)t/2.$$

If we switch to the new share, the expected post-tax proceeds after one year are

$$(1+n) - (1-b)t - ((1+n) - 1)t/2.$$

Equating these gives

$$1+n = (1+o) + (1-b)(t/2)/(1-t/2).$$

Table 7 shows figures for required capital return n on the new share for a range of basis costs and expected capital returns o on the old share, based on the expression above. No allowance has been made for discounting of tax payments (this could easily be done, but for practical purposes any increased accuracy is spurious).

¹³ This convention will be less realistic when the next taper threshold is close, because we would probably sell the original share soon after accrual of the threshold, rather than 12 months hence. But some assumption is needed to make any progress with the analysis. A possible alternative is to assume that both old and new shares are sold shortly after accrual of the next taper threshold on the old share. But this means that a high rate of tax is always paid on the new share, which seems a less realistic assumption on my experience.

¹⁴ It is simple to add discounting, but it makes the algebra more complicated, and the tables less easy to interpret; and for plausible rates, it does not make any difference to the general pattern of the results. Ignoring dividends can be justified either by the dividends actually being nil, or by assuming that the *difference* in the dividend streams for the new and old shares is small enough to ignore.

Table 7: Personal account: AIM share bought less than 1 year ago: capital return required on new share over period to one-year taper relief threshold to justify switching before the threshold

Basis cost of old share	Expected capital return on old share						
	-30%	-20%	-10%	0%	10%	20%	30%
0	-5%	5%	15%	25%	35%	45%	55%
0.2	-10%	0%	10%	20%	30%	40%	50%
0.4	-15%	-5%	5%	15%	25%	35%	45%
0.6	-20%	-10%	0%	10%	20%	30%	40%
0.8	-25%	-15%	-5%	5%	15%	25%	35%
1	-30%	-20%	-10%	0%	10%	20%	30%

Assumptions:

- Effective tax rate of 40% before the taper relief threshold, falling to 20% after
- Whatever share we are holding 12 months hence is sold at that time
- Tax losses assumed to be set against gains with the same taper as would have applied if the loss had generated a gain.

Some points of interest in Table 7 are as follows. For basis cost 0, across the first line of the table, we can see that the required *excess* capital return on the new share relative to the old is always 25% (the extra tax paid by switching is 0.2, so the incremental return required to compensate for this is 1/0.8). Moving down the table, basis cost increases (that is, accrued gain and tax cost of sale decreases), and the required excess capital return to justify a switch declines. For basis cost 1.0 (corresponding to an unchanged share price), there is no tax cost of switching, and so equal returns from old and new shares produce equivalent proceeds in one year.

(ii) *Original share bought between 1–2 years ago (so sale in 12 months of new share taxed at t)*

If the old share is retained, the post-tax proceeds are

$$(1+o) - ((1+o) - b) t/2$$

If the switch is made, the post-tax proceeds are $(1+n) - (1-b)t - ((1+n) - 1) t$

Equating these gives

$$(1+n) = ((1+o)(1-t/2) - b t/2)/(1-t).$$

Table 8 shows figures for required capital return n on the new share for a range of basis costs and expected capital returns o on the old share, based on the expression above. No allowance has been made for discounting of tax payments.

Table 8: Personal account: AIM share bought between 1 and 2 years ago: capital return required on new share over period to two-year taper relief threshold to justify switching before the threshold

Basis cost of old share	Expected capital return on old share						
	-30%	-20%	-10%	0%	10%	20%	30%
0	-21%	-10%	1%	13%	24%	35%	46%
0.2	-24%	-13%	-1%	10%	21%	33%	44%
0.4	-26%	-15%	-4%	8%	19%	30%	41%
0.6	-29%	-18%	-6%	5%	16%	28%	39%
0.8	-31%	-20%	-9%	2%	14%	25%	36%
1	-34%	-23%	-11%	0%	11%	23%	34%

Assumptions:

- Effective tax rate of 20% before taper relief threshold, falling to 10% after
- Whatever share we are holding 12 months hence is sold at that time
- Tax losses assumed to be set against gains with the same taper as would have applied if the loss had generated a gain.

The most important point to note in Table 8 is that the *differences* between return on the old share (top of each column) and equivalent-proceeds return on the new share (down each column) are smaller than in Table 7. This is as expected, since the tax rate t is lower.

More detailed observations on Table 8 are as follows. For basis cost 0, across the first line of the table, the required $(1+n)$ return on the new share is $(1-t/2)/(1-t)$, that is 1.125x the $(1+o)$ return on the old share. Moving down the table, basis cost increases (that is, accrued gain and tax sale cost decreases), and so the required return on the new share reduces.

In most of the left hand side of Table 8, the required capital return on the new share is negative. This is an artifact of one of our assumptions. We have assumed that any realized losses are set against gains with the same accrued taper as would have applied if the realized loss had been a gain. For example, if we switch and both the old share and the new share then fall, the loss realized on the new share is assumed to be set against a gain taxed at 20%; but if we had held on to the old share, the loss realized on the old share would have been set against a gain taxed at only 10%. The larger future loss relief outweighs the extra tax cost of switching now. However, in practice we might not always have gains taxable at 20% to set a loss on the new share against, and so this result should be treated with caution.

In any case, in practice we should always be able to switch into something with positive expected absolute return (if necessary, this can be cash). The negative figures everywhere to the left of the column headed 0% suggest the rule: when the original share is expected to fall, and the first year's taper relief has already accrued, ignore tax in decision making.

This point can be generalized, by noting the generally modest differences, even on the right side of Table 8, between the column head figures (returns on old share) and the figures in the body of the table (required returns on new share). This suggests a rough guideline: *once the first year's taper relief has accrued, ignore tax in decision making*. Another way of seeing this is to note that the first year threshold for taper relief represents the large marginal fall in tax (from 40% to 20%). The second year threshold is much smaller (from 20% to 10%).

8. Spread bets for short-term trades

A full discussion of the opportunities (and pitfalls) of spread betting would require another paper. However, this present paper would appear incomplete to many investors if it did not note that in the UK, spread betting offers significant tax advantages when the investor wishes to take a position which is expected to be held only a few weeks or months. Spread bets avoid stamp duty of 0.5% on share purchases. Spread bet winnings are not taxable, and correspondingly there is no relief for losses. The bookmaker pays a tax of 3% of gross profits, that is 3% of the amounts lost by clients. Bets are available on many liquid assets including a wide range of quoted shares, and at some bookmakers this extends down to companies with a market capitalization as small as £10m. The example below considers a bet on an individual quoted share.

A spread bookmaker quotes to its clients (through an internet platform, or on request over the telephone) a bid-ask spread for what is essentially a futures contract on a share price. By custom (adopted by all spread betting firms) the futures contract that is the bet expires quarterly, on the Tuesday prior to the third Wednesday on any of the four quarter months (March, June, September or December). The futures price incorporates a funding cost (typically LIBOR +1-2%), and adjustment for any dividends expected to be due over the term to expiry of the bet.

Although spread betting is a retail product, positions giving economic exposure running into many millions of pounds can be taken, subject to the client's credit standing, and the bookmaker's ability to hedge in the underlying market. For modest sizes (typically up to Normal Market Size or sometimes a little larger), quotes are available to execute online or immediately over the telephone; the bookmaker does not hedge every trade in popular shares, but offsets client buys and sells where he can, and then hedges if his net exposure exceeds a limit. For orders in larger sizes or more obscure shares, the bookmaker will generally seek to hedge the entire order in the market, and fill the client's order on a piecemeal basis as and when the hedge can be executed.

Example

Suppose it is now 1 May, LIBOR is 4¾%pa, and the bookmaker's interest margin is 1¼%pa. The bid-ask spread for a share is 70-72p in the market and the share is expected to go ex-div a dividend of 2p on 14 May. The spread bet middle price for the September contract will be approximately

$$71p \times (1.06)^{199/365} - 2p = 73.4p$$

where interest on the dividend has been ignored (as I believe is common). The bookmaker applies market spread and an additional bookmaker's spread of say 0.3% either side around this price, and makes a September bid-ask quote of 72.2-74.6 to the client. Bets are expressed in £ per penny on this price. For example, a client who wants an economic exposure equivalent to owning 10,000 shares would "buy £100 of September at 74.6."

After the bet has been opened, the client's position is marked to market in real time. A client with a position of size £100 gains or loses £100 for every penny that the price moves. If a large price movement against the client occurs, the client may be required to pay a margin call, that is deposit funds with the bookmaker to cover the marked-to-market losses.

When the client wishes to close the position, she gives an equivalent but opposite order to the original opening order. In our example above, suppose that 30 days later the quote for the share in the market is now 80-83p, and the share has now gone ex-dividend. The spread bet middle price for September will be approximately

$$81\frac{1}{2}\text{p} \times (1.06)^{169/365} = 83.7\text{p}$$

After applying market spread and the bookmaker's spread (assumed at 0.3% either side), the September bid-ask quote would be say 82.0-85.5. To close her position, the client would "sell £100 of September at 82.0." The client's profit on the trade is $(82.0-74.6) \times £100 = £740$.

Bookmakers will typically quote for at least the next three quarterly expiries, that is up to nine months ahead. Although it is uncommon, positions can in practice be maintained for longer than this, because bets which are close to expiry can be "rolled over" into the next contract. On a rollover the funding cost to the next expiry plus a bookmaker's spread will be charged; but typically the spread would be charged on only "one side" (eg on opening of the new contract, but not on the closing of the expiring one), and at a reduced rate compared to the normal spread for a new trade.

Comments on spread betting

A disadvantage of spread bets is that the bookmaker's spread tends to be slightly wider than commission one might pay to a stockbroker. This partly reflects the gross profits tax which bookmakers pay of 3% of the amounts lost by clients. A more important disadvantage is that a funding cost (typically at around LIBOR +1-2%) is charged on the entire position (usually this applies even where the client has been required to deposit some funds as margin with the bookmaker). I suspect that a large part of the bookmakers' income is attributable to margins in the funding charges. For long positions held for more than a few weeks, the cost of funding tends to outweigh saving on stamp duty (although if the investor wants leverage, a spread bet is still probably cheaper than most alternatives).

Spread bets may be relatively more attractive as a mechanism for holding short positions rather than long positions. This is partly because it is often difficult for a personal investor to take a short position by any other means¹⁵. Another reason is that the two-way forward price quoted by a spread bookmaker on the internet or in response to a client enquiry generally seems to be a symmetrical spread around the current share price projected to the bet's expiry date at LIBOR +1% or 2% (rather than, as might be expected, a margin *over* LIBOR for the customer's buy price, and a margin *under* LIBOR for the customer's sell price). I believe this practice reflects the fact that the spread bookmakers' clients in aggregate tend to have a net long position. But for the customer who uses the spread bookmakers to take a short position, it means receiving the benefit of LIBOR +1% or 2% tax-free. This is more favourable than a contract for differences, where the expected asymmetry prevails: a financing cost *above* LIBOR for funding a long position, and a financing credit *below* LIBOR representing interest on the sale proceeds of a short position.

¹⁵ Except a contract for differences (CFD), which is broadly similar to a spread bet, but with two main differences: a CFD is open-ended with no specified expiry date, and gains and losses on a CFD are chargeable to CGT. The interest charged to reflect funding of a long position (or credited on a short position) is effectively included in the charge to CGT; the same applies to net dividends credited on a long position (or debited on a short position).

It is difficult to make money in the long term from short selling. The general drift of share prices is upwards; and a short position which goes against the investor can quickly become a very large problem. Short selling may not be a good idea; but if one is going to try it, one might as well choose the least disadvantageous mechanism available. If the investor understands the perils of excessive leverage *and can avoid the temptations in practice*, a spread bet may be the most cost-effective way of making a short-term trade, particularly a short position.

A detour: compounding, volatility and leverage

Although spread betting seems theoretically attractive for short-term trades, it is widely acknowledged that a majority of spread betting clients lose money over time. A losing proportion of 85% to 90% of clients is commonly quoted, and there seems no reason why this should be understated. The disposition of the tax authorities, whereby winnings are tax-free and bookmakers pay a tax of 3% of the amounts lost by clients, is consistent with this impression. The losses are partly explained by the bookmakers' spreads and funding charges. But another explanation may be the substantial leverage which is available – for single shares, anything from 2x to 10x the margin required from the investor. Even if this leverage were free (which it is not), increasing leverage beyond a certain level reduces expected compound return. This is very important to the investor and not always widely understood, so it is worth a detour to give some mathematical insight.¹⁶

From a mathematical perspective, the effects of leverage and volatility and higher moments on compound return can best be appreciated by considering the Taylor series expansion of the expected log return about the expected return. Define the simple arithmetic return r as the ratio of the end price to the start price. Writing $E[r] = E$, the expected log return can be found from the first terms of the Taylor series expansion of $\log(1+r)$ about the mean return E as follows

$$f(r) = f(E) + (r - E)f'(E) + (r - E)^2 f''(E) / 2! + (r - E)^3 f'''(E) / 3! + \dots \text{ (Taylor's theorem)}$$

which for the function $f(r) = \log(1+r)$ gives

$$\log(1 + r) = \log(1 + E) + (r - E)/(1 + E) - (r - E)^2/(2(1 + E)^2) + (r - E)^3/(3(1 + E)^3) - \dots$$

For this series to converge when we take expectations, we need $E[(r - E)^n] < (1 + E)^n$. We can usually meet this requirement by taking a time period sufficiently short. Then taking expectations of both sides, the second term on the right is the first moment about the mean, which is zero; and the numerator in the third term is the variance of r , say V . The fourth and fifth terms can be related to the skewness of r , say S , which is $E[(r - E)^3]/V^{3/2}$; and the kurtosis of r , say K , which is $E[(r - E)^4]/V^2$. So we can write

$$E[\ln(1 + r)] = \log(1 + E) - V/(2(1 + E)^2) + S V^{3/2}/(3(1 + E)^3) - K V^2/(4(1 + E)^4) + \dots$$

This expression gives an immediate insight into how variance of return reduces compounding, as do negative skewness and high kurtosis.

Now suppose that the return r is leveraged by L , where $L > 0$ represents a multiple of the investor's equity available for investment¹⁷. The expected log return on the investor's equity becomes

¹⁶ For alternative expositions of this material, see Booth & Fama (1992) and Wilcox (2004).

¹⁷ $0 < L < 1$ corresponds to investing only part of one's equity, leaving the rest in a risk-free asset; $L > 1$ corresponds to borrowing.

$$E[\ln(1+Lr)] = \log(1+LE) - L^2V/(2(1+LE)^2) + L^3S V^{3/2}/3(1+LE)^3 - L^4K V^2/4(1+LE)^4 + \dots$$

The negative second term containing L^2 and V means that the expected log return on a leveraged investment with high volatility can easily become negative, within the range of leverage levels (eg $2 < L < 10$) commonly provided in spread betting, or indeed by prime brokers to some hedge funds. Expected log return can easily be negative, *despite the expected return E being positive*.¹⁸ It is expected log return, *not* expected return, which is the better indication of the results an investor can reasonably expect from long-term compounding.¹⁹

For negatively skewed or fat-tailed returns, the higher terms in the Taylor series expansion involving L^3 and L^4 can become important. This gives the insight that *the higher moments of return matter more to a leveraged investor*.

Looking up logarithms is inconvenient for mental arithmetic, so for this purpose it is helpful to have a simpler approximation. For short periods E will typically be small, so LE will be small, and $\log(1+LE) \approx LE$, and $(1+LE)^2 \approx 1$; and if we use these approximations for the first and second terms of the expansion, the errors in the two approximations are of opposite sign. A quick mental arithmetic approximation of the expected log return of a leveraged investor is then

$$LE - L^2V/2$$

and it is often useful to keep this expression in mind.

The approximation may be inaccurate if the return has a pathological distribution with unusual higher moments, but it is often useful as a quick first check of whether a proposed level of leverage is likely to be sensible. Wilcox (2004) uses this approach to give good quick insights into a number of failed financial structures, including leveraged hedge funds, constant proportion portfolio insurance, and split capital investment trusts.

¹⁸ To see this for a trivial case, consider an investment which either doubles or falls 60% with equal probability in every period. Expected return is $(2+0.4)/2 = 1.2$, a return of 20% per period. But expected log return is $(\log(2)+\log(0.4))/2 = -0.11$, which gives a median compound return of 0.89, that is -11% per period, so in the long term you almost surely go broke despite a positive expected return. (The $E-V/2$ approximation for expected log return gives $0.2 - 0.64/2 = -0.12$, which is plenty good enough to show that trouble is likely.)

¹⁹ The distribution of terminal wealth from any long period of compounding is likely to be very positively skew – a small number of very high outcomes contribute a lot to the mean. For this reason median terminal wealth – found by summing the log returns, and then taking the anti-logarithm – is a much better measure of typical results. This can be illustrated using the same example as in the previous footnote, with a time horizon of 3 periods. This gives eight equiprobable outcomes for terminal wealth levels: 0.064, 0.32 (three times), 1.6 (three times), 8. Taking the mean of these, we find mean terminal wealth as 1.728 (=1.2³). But 7 of 8 possible outcomes are below this. The median terminal wealth of $\exp\{-0.11 \times 3\} = 0.72$ seems a better measure of typical results.

9. An attempted synthesis: practical guidelines for the active personal investor

This section attempts to draw together earlier sections of the paper to provide some practical guidelines for the active personal investor.

Top-level principles and concepts

- *Post-tax thinking on risk and return.* Risk and return should always be considered on a post-tax basis. The taxation authority acts as a profit-sharing partner who shares in the investor's profits and losses, but contributes no capital. The presence of this partner probably means that the investor should be more tolerant of downside risk in taxable accounts than in tax-advantaged accounts.
- *Capital gains taxes as transaction costs.* Decisions to vary a portfolio should be considered by opportunity cost comparisons, including tax effects. Income taxes can be regarded as reductions in yield in perpetuity, which are relatively simple to model. Capital gains taxes are more difficult, and probably best regarded as (sometimes very large) transaction costs.
- *Judicious lethargy.* Tax transaction costs are often high and always certain. The pre-tax benefits derived from a trade are usually uncertain; and most investors tend to be overconfident about them, most of the time. Hence trades should be made only when one is very sure that the existing portfolio could be improved. Judicious lethargy can be a valuable trait.
- *Loss harvesting.* An exception to the previous point is where tax transaction costs are negative, that is a share in a taxable account is standing at a loss and realizing the loss will enable it to be offset against gains. In these circumstances, one should realize the loss more readily than in a tax-free account.
- *A unified view across all taxable and tax-advantaged accounts.* Tax-aware management is best achieved by taking a unified view of decisions across all accounts held by one investor (and perhaps one family). For example, when a new share is purchased, one needs to consider whether it is of a type best allocated to ISA, SIPP or a taxable account (some guidelines for this are given below). When reducing a shareholding standing at a gain held in more than one account is being considered, one would normally want to sell in a tax-advantaged account first – unless there is a requirement to create a gain in the taxable account to utilise the annual exemption, or to offset other losses. Making such decisions in a coherent way requires a unified view across all the accounts.

Allocating new purchases to ISA, SIPP, corporate and personal accounts

When a new purchase of shares is made, the investor needs to decide to which account to allocate it. Table 9 summarises appropriate preferences for an investor who has a spread betting account, an ISA, a SIPP, a company and a personal account. I explain these preferences in more detail below, by reference to the four boxes in the body of Table 9.

Table 9: preferences for allocating new purchases of different types of share (PA = personal account)

Listing of stock	Expected holding period	
	Short (say <1yr)	Long (say >1yr)
Fully listed (or non-trading companies on AIM)	1. Spread bet (if very short) 2. ISA 3. SIPP 4. Corporate 5. PA	1. Corporate 2. ISA 3. SIPP 4. PA
Trading company listed on AIM	1. SIPP 2. PA 3. Company	1. SIPP ≈ PA 2. Company

Upper left box of table. If a view on a fully-listed share is expected to be held only for a very short time, or a short position is desired, then a spread bet should be considered. In all other cases, a fully-listed share should be allocated to a tax-deferred account in preference to personal account; and to the ISA in preference to the SIPP, reflecting the fact that the ISA *must* be filled with fully-listed shares, whereas the SIPP can hold AIM-listed shares as well. If neither the ISA nor the SIPP have spare capacity, allocating to the investor's company is probably preferable to personal account, for two reasons. The first reason is that any dividend income will be received without further tax in a company, but would be taxed at 25% in the personal account. The second reason is that in personal account there is a *relative* advantage (faster taper relief) in holding trading companies listed on AIM rather than fully listed shares. Hence one should generally not fill up the personal account with fully listed shares if there is spare capacity in an alternative location.

Lower left box. Trading companies listed on AIM cannot be held in the ISA. They are tax-free in a SIPP, and so this is the first preference. If there is no spare capacity in a SIPP, then allocating them to personal account gives full taper relief after two years; this is preferable to a company, where there is no taper relief.

Upper right box. If a fully-listed stock (particularly one with a high dividend yield) is expected to be held for a long period, this is probably the most suitable type of investment to allocate to the investor's company. Any dividends are received tax-free, *and* the ex-dividend adjustment in the market (which invariably seems to be on a gross basis) reduces the share price a little every six months, reducing unrealised CGT. However these stocks can also be allocated to the ISA or SIPP (in the same order of preference and for the same reasons as on the left hand side of the table). Allocating them to personal account is the last resort.

Lower right box. Trading companies listed on AIM cannot be held in the ISA. They are tax-free in the SIPP, but if held for a long period are also lightly taxed in personal account (just 10% CGT after two years of taper). In the personal account, trading companies listed on AIM enjoy a large *relative* advantage relative to a fully-listed share, but there is no such relative advantage in a SIPP. Taking all these considerations into account, there seems to be little to choose between allocating trading companies listed on AIM to a SIPP or a personal account. In a corporate account, there is no taper relief on holding AIM shares long-term, only indexation relief. This will generally be much less valuable than accelerated taper relief, and so such shares should generally not be allocated to the company.

Other guidelines on allocating shares to accounts

Some of the following guidelines are implicit in the table above, but it may be helpful to state them separately.

- Fully listed companies should generally be held in an ISA, or else in a SIPP; or in a corporate account, if a long holding period is expected.
- If, exceptionally, fully listed companies are held in a personal account, they should be chosen with a view to holding for a *very* long time (maximum taper relief does not accrue until 10 years after purchase).
- For any positions which are expected to be held only for a short time (a few weeks or months), consider spread bets. Short positions should generally be taken as spread bets. But be very careful about leverage. Remember that higher moments can be a drag on compound returns, and that leverage can increase the drag: remember the expression $LE - L^2V/2$.
- For the investor's personal account, a portfolio focused on trading companies listed on AIM offers substantial tax advantages. It may well be the only approach under which skilled active management of a personal account has some realistic prospect of producing post-tax excess capital return.
- It should be rare for fully listed companies to end up in the investor's personal account, because there are several more tax-efficient accounts in which to hold them (ISA, SIPP or corporate account).

Further guidelines for personal account

The following further guidelines are appropriate for management of a personal account comprised of trading companies listed on AIM:

- The investor should try to hold each share for at least one year, because this gives a large marginal reduction in tax (from 40% to 20%).
- After a share has been held for one year, tax becomes a much less important consideration: the difference between 20% and 10% marginal tax is not very significant. As a rough guideline: *once an AIM share has been held for one year, ignore tax in decision making.*
- The investor should probably try to avoid adding to a holding in a trading company listed on AIM more than 3 months (say) after first purchase, because this delays the option to reduce the holding with some taper relief. (This is explained further in the “quirks” section below.)

Further guidelines for corporate accounts

The following further guidelines are appropriate for management of a corporate account:

- A corporate account, viewed in isolation, is indifferent between trading companies listed on AIM and fully listed shares. (This is because there is no taper relief for companies, and so the two types of shares receive the same tax treatment in a corporate account.)
- Nevertheless, it should be rare for shares in trading companies listed on AIM to end up in the investor's corporate account. (This is because there are two better choices for these shares: they are the preferred holdings for personal accounts, and can also be held in a SIPP.)
- There may nevertheless occasionally be good reasons for allocating AIM shares to a corporate account. In particular, possible reasons are to avoid deferring taper relief on a previous personal account investment in the same company's shares; or to avoid a very high

dividend being taxed at 25% in personal account. These ideas are discussed further under the next two headings below.

In a personal account, increasing (reducing) existing holdings may delay (advance) taper relief

In a tax-advantaged account or a corporate account, a decision to add to an existing shareholding can be considered on a purely contemporaneous basis, by reference to the current portfolio, the current share price, and alternative investment opportunities. Details of the dates, prices and amounts of previous purchases of the same share are irrelevant to the decision. Similarly, a decision to reduce an existing shareholding can also be considered on a purely contemporaneous basis, by reference to the current share price, the current holding and its *average* basis cost, and hence the taxable gain or allowable loss (if any) crystallised by the marginal sale. However, for personal accounts the decision is more complicated, because taper relief depends on the full sequence of purchases in any share. This is particularly important for trading companies listed on AIM, where taper relief accrues rapidly. The time-dependent nature of taper relief may make it *less* attractive to make marginal *additions* to holdings of a particular share in a personal account more than three months (say) after one's first purchase (see example 1 below). Conversely, a long history of earlier purchases may make it *more* attractive to make marginal *reductions* to an existing holding (see example 2 below).

Example 1

Suppose a share in a trading company listed on AIM was first bought 11 months ago, and has risen at a rate we consider fair in the light of news since then. We now have new cash to invest, and the share which we first bought 11 months ago is again available at a price which would be attractive for purchase in a tax-advantaged account. However, we are considering purchase in a taxable personal account. Suppose that the new purchase for a personal account is *not* made. Then one month from now, the "old" share could be sold with 50% taper relief, that is the contingent tax on any gain on the "old" share will fall in one month from 40% to 20%. So if the price rises somewhat faster than we consider justified in the next few months, we might well find it attractive to reduce the holding.

Now suppose, on the other hand, that the new purchase of the shares after 11 months *is* made. Then the Last In, First Out (LIFO) identification rule means that subsequent sales will generally be matched against the latest purchase. The new purchase will therefore defer by 12 months the *option* to sell the earlier purchase subject to a tax rate of only 20%. In the light of this, it might be better to invest in some other share with broadly similar prospects, rather than make the marginal addition to an existing holding, which defers the *option* to realise a gain subject to only 20% tax.

Example 2

Suppose we disregarded our guideline in Example 1 (perhaps the share was irresistibly attractive!), and made a marginal addition to the existing holding 11 months after first purchase. After another 6 months, the share price has risen considerably more than we consider is justified by news. If we now make a marginal reduction in the holding, then under the LIFO rule the marginal sale will crystallise CGT at 40% of the gain. If we viewed the decision on a purely marginal basis, this 40% penalty might well deter us from selling the share, even though we consider it somewhat overvalued. But the decision to sell in a personal account should not be made solely by reference to tax crystallised by this marginal sale. We also need to consider the advancement of the *option* to make further sales of the earlier purchases, at a lower tax penalty.

That is, once we have sold the shares we bought 6 months ago subject to 40% CGT, we have created the option to make further sales subject to only 20% CGT.

The general principle from these examples is that in a personal account, decisions to add to or reduce existing holdings cannot properly be made on a purely contemporaneous basis. One also needs to consider that marginal purchases (sales) can defer (advance) the option to sell earlier purchases at lower tax penalties. Details of the full sequence of past purchases are relevant. Fortuitously, applying this principle also tends to promote diversification.

In practical management of a sizeable portfolio, a holding in a particular share may be increased over several weeks, or even months (notwithstanding the point in Example 1). There might be a dozen or more dates and quantities of purchases, and hence a dozen or more anniversary dates on which an extra year of taper relief accrues for some part of the holding (if not already eliminated under the LIFO rule). In day-to-day management of a portfolio, it helps to have a system for monitoring and highlighting the sequence of such anniversary dates for every share.

High dividend yields best allocated to corporate account

Shares with a high dividend yield are probably most appropriately located in a corporate account. A dividend from shares in a UK company is received as franked investment income, which means there is no tax to pay. There is a second benefit in that the ex-dividend adjustment when the share goes ex-dividend reduces the prospective capital gain which will crystallise on any future sale of the share. In contrast, in a personal account a dividend will be taxed at 25% of the cash amount received by the investor. In an ISA or a SIPP, a high dividend yield carries no tax advantages or disadvantages.

10. Some further quirks

A personal investor may find it useful to be aware of the following further quirks in the capital gains tax (CGT) legislation.

Using tax sale costs when implementing top-down asset class re-weightings

Most of our discussion so far has focused on a scenario where the investor is considering a marginal switch from one overvalued asset to another cheaper asset. Sometimes the investor's motivation for selling may be more of a top-down view: for example when the investor is bearish about equities as an asset class, he may want to switch part of his portfolio to cash, without having any views about which particular shares should be sold. In this situation, the shares selected to be sold should first be those in tax-advantaged accounts, and then those with the lowest tax sale costs. In practice, this may mean that top-down asset class re-weightings are implemented mainly via the tax-advantaged accounts. The same applies when making other types of top-down reallocations, for example between sectors or countries.

You can sell and buy back within 30 days without long-term CGT implications

Prior to 6 April 1998 it was customary for investors to ensure full utilisation of the annual CGT exemption by “bed-and-breakfast” transactions, whereby shares were sold near the close of trading one day and re-purchased early the next day. New matching rules introduced with effect from 6 April 1998 require sale to be matched against a repurchase of the same shares within the following 30 days. This “anti bed-and-breakfast” rule effectively means that one can sell and buy

back shares within 30 days without long-term CGT implications. After the sale and buy-back, the original base cost of shares purchased many years ago is restored.

One consequence of this rule is that it may be best to sell *immediately* if one learns bad news (for example, a profit warning), and then make a more considered decision about what one wishes to do in the next 30 days. If one decides to repurchase, and the price has fallen a little further in the 30 days (that is, the decision to sell immediately turns out to have been correct), you get a small gain; if the price actually rises in to 30 days (that is, the decision to sell immediately turns out to have been wrong), you get a small loss. Either way, the 30-day rule means that an *immediate* decision to sell on bad news can be made, without necessarily crystallizing a large long-term gain. This scenario is an exception to the principle of judicious lethargy.

Example

Suppose the investor bought shares in a fully-listed company just under 3 years ago at 100p. The share price at close of trading yesterday was 600p. A profit warning is released to the market at 1100 hrs today. Between 1101 and 1115 hrs, the investor sells all her shares into a falling market, at an average price of 500p.

Over the next two weeks the investor investigates the circumstances which led to the profit warning, and forms the view that the price fall is over-done. She therefore re-purchases her shares at an average price of 480p.

The earlier sale is matched against the purchase within 30 days to give a gain of $(500 - 480) = 20p$, giving rise to CGT of 8p. The investor's base cost for holding is now 100p, and the accrued holding period for taper relief purposes is 3 years. The decision to sell *immediately* on reading the profit warning has crystallised tax of just 8p (plus transaction costs). Were it not for the 30-day rule, selling *immediately* on reading the profit warning would have been a much more difficult decision, because the gain crystallised by the immediate reaction (which might well be thought to be wrong a few hours later) would have been 400p, giving rise to tax of 160p.

Whilst this point is worth knowing, in practice I think it is probably useful mainly for an investor with a very small holding, which can be sold quickly without large market impact.

Where a share was sold in the month from 6 March to 5 April, repurchase and resale can shift CGT forward one year

If a share has been sold at a large gain in the month from 6 March to 5 April, it may be worth deferring most of the CGT for one year by repurchasing and then reselling, within 30 days of the original sale but after 5 April. (For example, under the 30-day rule the March 2008 sale is matched against the April 2008 buy, giving only a small gain or loss for 2007/08, based on the price change over one month; then, the next year, the April 2008 sale will be matched against the original purchase of the stock (eg in 1999 say), so that the large gain crystallises in the 2008/09 CGT return.)

This technique can also be useful in certain circumstances where a share has been sold at a *loss* in the period from 6 March to 5 April. If the loss reduces the net chargeable gains for the year below the annual exemption, part or all of the exemption will be wasted. This waste of the exemption can be avoided by repurchasing and then reselling an appropriate number of shares within 30 days, but after 5 April as above.

Using CFD's to defer gains by one year

A similar technique which can be used at any time during the year (rather than just in the month leading up to 5 April) is to open a contract for differences (CFD) which offsets the original position. A contract for differences is broadly similar to the spread bet described earlier, except that gains or losses are taxable. The contract is made between the investor and the CFD provider (this may be an investment bank, or alternatively a firm which also offers spread betting). The contract specifies a number of shares and an opening price, and provides for the *difference* between this value and the corresponding value at the time when the investor closes the contract to pass between the investor and the CFD provider. In contrast to the limited life of spread bet, a CFD usually does not specify a termination date. While the position remains open, the investor will be charged a daily financing cost (typically LIBOR +1-2%) to fund a long position, or receive a financing credit on short positions (typically LIBOR – 2%). The CFD provider hedges the position in the cash market in the same manner as a spread betting firm. No stamp duty is payable on a CFD.

For tax purposes, this technique can defer a gain to the subsequent tax year, when the annual exemption or other losses may be available to offset it. Similarly a loss can be deferred if appropriate, or the holding period of the original shares for taper relief purposes can be extended. The essential point is that the contract for difference is a distinct asset from the original shares.

Example

Suppose 10,000 shares were originally purchased at 100p, and on 6 January 2008 stand at 200p. The investor wishes to lock in the economic gain, while deferring the chargeable gain until after 5 April 2008. The investor takes a short position via a contract for differences over 10,000 shares at 200p.

On 6 April 2008, the share price is 250p. The investor sells the 10,000 shares, and closes the short position on her contract for differences over 10,000 shares. Ignoring brokerage charges and financing credits on the CFD short position, her gains for the tax year commencing on 6 April 2008 are

Gain on shares	15,000
Loss on CFD	<u>(5,000)</u>
Net chargeable gain	10,000

The £10,000 net chargeable gain is the same as if she had sold the shares in January 2008, except that it has now been realized in the 2008/09 tax year.

Interaction of losses, taper relief, and the annual exemption

The interaction of current year losses, taper relief and the annual exemption detailed in section 4 above means that current year losses sometimes lead to “waste” of both taper relief and annual exemption. The interaction is slightly more favourable for unused losses brought forward from prior years – these sometimes lead to “waste” of taper relief, but not the annual exemption.

Example

Relevant information for the year ended 5 April 2008 is as follows:

	£
Gain realised in 2006/07 on share A (eligible for 50% taper)	20,000
Potential taper relief on gain on share A	15,000
Loss realised in 2006/07 on share B	16,000
Loss on share C brought forward from a prior year	50,000
Annual exemption	9,200

The investor's net chargeable gains for the year ended 5 April 2008 are calculated as 20,000 – 16,000 = £4,000. After taper relief, this reduces to £2,000. The annual exemption of £9,200 is applied against this gain, giving a CGT liability of Nil. The taper relief on share A is “wasted”, and so is the unused £7,200 of the annual exemption. But the prior year loss of £50,000 on share C can be carried forward again.

In the absence of the current year loss of £16,000 on share B, the gain on share A would have been reduced by the loss brought forward on share C, but only to the level of the annual exemption (ie 20,000 – 10,800 = 9,200, with the other 39,800 of the loss on share C carried forward again). That is, the taper relief entitlement on share A would have been “wasted”, but the annual exemption would not have been.

Setting losses against gains with varying accrued taper reliefs

Where the investor has losses, together with different gains with different accrued taper reliefs, the legislation provides that losses are to be set against gains in the order which gives the largest reduction in chargeable gains. In other words, *losses are first set against gains with the least taper relief.*

For this purpose it is worth noting that a single sale transaction can sometimes give rise to *two* gains, with different rates of taper relief. This applies where the investment's “business asset” status changes during the holding period. For example, if a trading company listed on AIM moves to the Full List, it will be a business asset for the period listed on AIM and a non-business asset thereafter. The overall gain is apportioned into two gains by reference to the periods as a business and non-business asset. The rate of taper for each gain is based on the overall combined holding period. This is illustrated in the following example.

Example

Suppose the investor buys £100 of shares in a trading company listed on AIM on 1 January 2007. On 1 January 2008 the company moves to the Full List. On 1 January 2012 the investor sells the shares for £200.

The overall gain of £100 is apportioned based on 1 year of business asset ownership and 4 years of non-business asset ownership:

Business asset gain : 1/5 x £100	=	£20
Non-business asset gain: 4/5 x £100	=	£80.

The appropriate rates of taper for each of these gains is determined by reference to the overall period of ownership, ie 5 complete years. This gives 75% taper on a business asset, and 15% taper on a non-business asset. So the tapered gains are £5 and £68 respectively.

If the investor has losses to set against the gains, they are set against the non-business asset (which will receive very little taper relief), rather than the business asset (which will receive 75% relief anyway).

11. Conclusions

This paper has reviewed the interaction of United Kingdom taxation and portfolio decisions by a personal investor managing her own investments in quoted company shares. Practical guidelines were given in section 9 and 10. The most important top-level principles and concepts were:

- think post-tax, not pre-tax, when evaluating risk and return;
- when considering possible changes to a portfolio, allow for capital gains tax, which should be treated as (large) transaction costs;
- judicious lethargy (because the tax transaction costs of varying a portfolio are often high and always certain, whereas any advantage gained from the variation is uncertain);
- loss harvesting (an exception to judicious lethargy – where tax transaction costs are negative);
- in personal accounts, decisions to add to or reduce existing holdings are not always purely contemporaneous – they can modify taper relief options on earlier purchases as well;
- take a unified view across all the investor’s taxable and tax-advantaged accounts.

This paper has taken the perspective of an active personal investor managing her own investments. Many investors delegate portfolio management to professional managers. Assuming the professional management is active rather than passive, many of the concepts in this paper seem applicable. It would be interesting to see papers on tax-aware portfolio management written from the perspective of professional managers in the UK, rather than the US.

Many of the concepts I discuss can also be applied, with suitable adjustments, to management of the portfolios of other entities which pay tax on their investments – for example, insurance companies, other companies, and most private trusts. It would be interesting to see papers written by people who do this.

Finally, although it is beyond the scope of this paper, I note that the most tax-efficient thing either a UK-resident individual or company can do with shares standing at a large gain is probably to give them to charity. Generally, the entire gain will be exempt from tax, *and* the market value of the shares gifted can be offset against other income. This often means that well over half the value received by the charity can come from tax relief.

References

- Adams, A.T. & Clunie, J.B. (2006) ‘Risk assessment techniques for split capital investment trusts’ *Annals of Actuarial Science*, **1**, 1: 7-36.
- Amromin, G. (2003) ‘Household portfolio choices in taxable and tax-deferred accounts: another puzzle?’ *European Finance Review*, **7**: 547-582.
- Arnott, R.D., Berkin, A.L & Ye, J. (2001a) ‘The management and mismanagement of taxable assets’ *Journal of Investing*, Spring, 15-21.

- Arnott, R.D., Berkin, A.L. & Ye, J. (2001b) 'Loss harvesting: what's it worth to the taxable investor?' *Journal of Wealth Management*, **3**, 4: 10-18.
- Berkin, A.L. & Ye, J. (2003) 'Tax management, loss harvesting and FIFO accounting' *Financial Analyst's Journal*, **59**, 4: 91-102.
- Bicksler, J.L. and Chen, A.H. (1985) 'The integration of insurance and taxes in corporate pensions strategy' *Journal of Finance*, **40**, 3: 943-955.
- Black, F. (1980) 'The tax consequences of long-run pension policy' *Financial Analyst's Journal*, **36**: 21-28.
- Blake, D. and Khorasane, M.Z. (2005) 'Commentary on the paper "Pension funds and capital structure: why hold equities in the pension fund?"' *North American Actuarial Journal*, October 2005.
- Booth, D.G. and Fama, E.F. (1992) 'Diversification returns and asset contributions' *Financial Analyst's Journal*, **48**, 3: 26-32,
- Constanides, G.M. (1983) 'Capital market equilibrium with personal tax' *Econometrica*, **51**: 611-636.
- Constanides, G.M. (1984) 'Optimal stock trading with personal taxes' *Journal of Financial Economics*, **13**: 65-89.
- Dammon, R.M., Spatt, C.S. and Zhang, H.H. (2004) 'Optimal asset location and allocation with taxable and tax-deferred investing' *Journal of Finance*, **59**: 999-1037.
- DeMiguel, V. & Uppal, R. (2005) 'Portfolio management with the exact tax basis via non-linear programming' *Management Science*, **51**, 2: 277-290.
- Garlappi, L. and Huang, J. (2006) 'Are stocks desirable in tax-deferred accounts?' *Journal of Public Economics*, **90**: 2257-2283.
- Jeffrey, R.H. (2001) 'Tax-efficient investing is easier said than done' *Journal of Wealth Management*, **3**, 2: 9-15.
- Jeffrey, R.H. and Arnott, R. (1993) 'Is your alpha big enough to cover its taxes?' *Journal of Portfolio Management*, Spring 1993.
- Osorio, M.A., Gulpinar, N., Rustem, B. & Settergren, R. (2004a) 'Post-tax optimisation with stochastic programming' *European Journal of Operational Research*, **157**: 152-168.
- Osorio, M.A., Gulpinar, N., Rustem, B. & Settergren, R. (2004b) 'Tax impact on multi-stage mean-variance portfolio optimisation' *International Transactions in Operational Research*, **11**: 535-554.
- Poterba, J.M., Shoven, J.B. & Sialm, C. (2000) 'Asset location for retirement savers' *NBER Working Paper No. 7991*.
- Ralfe, J., Speed, C. and Palin, J. (2004) 'Pension funds and capital structure: why hold equities in the pension fund?' *North American Actuarial Journal*, July 2004.
- Sherris, M. (1992) 'Reserving for capital gains tax (an application of option pricing theory)' *Journal of the Institute of Actuaries*, **119**, 1: 45-68.
- Shoven, J. B. and Sialm, C. (2003) 'Asset location in tax-deferred and conventional savings accounts' *Journal of Public Economics*, **88**: 23-38.
- Tepper, I. (1981) 'Taxation and corporate pension policy', *Journal of Finance*, **36**: 1-13.
- Wilcox, J. (2004) 'Survival of the fittest' *Journal of Asset Management*, **5**, 1: 13-24.
- Wilkie, A.D. (1995) 'More on a stochastic asset model for actuarial use' *British Actuarial Journal*, **5**: 777-964.
- Wilkie, A.D., Waters, H.R. & Wang, S.Y. (2003) 'Reserving, pricing and hedging for guaranteed annuity options' *British Actuarial Journal*, **41**: 263-425.

APPENDIX A

This appendix gives details some of the further rules relevant to self-invested pension plans (SIPPs).

The investor (and her employer) can make any level of contribution to any pension plan, and these contributions will generally qualify for income tax relief (if made by the investor) or corporation tax relief (if made by the employer), subject to two main limits on the relief. First, the investor's aggregate tax-relieved contributions to all pension plans of which he is a member are limited to the larger of £3,600 or 100% of employment income. Second, if the investor's annual increase in pension savings exceeds a limit (expected to rise to £255,000 in 2010/11), the excess attracts a tax charge of 40%. "Increase in pension savings" includes all contributions paid to registered pension schemes during the year by or on behalf of the investor, plus the increase in value (on a prescribed basis) of any benefits in defined benefit pension plans. So for an investor who has no defined benefit pensions, the effective limit for tax-relieved contributions by investor and her employer combined is expected to be £255,000 per annum (from 2010/11).

No income tax or capital gains tax is payable on investments held in the pension fund wrapper. So as with ISAs, SIPPs provide the twin advantages of tax -free roll-up and the ability to vary a portfolio without tax consequences. At any time after the investor's 55th birthday, pension income can commence; at the same time, one third of the fund used to provide a pension (that is, one quarter of the total fund utilised at the time) can be taken as tax-free cash (subject to a cumulative maximum of the one quarter of the "lifetime limit" detailed in section 2). Pension income can be drawn either by purchasing an annuity, or by "income drawdown" of the accumulated fund. Any rate of income drawdown can be selected, between zero and an upper limit of 120% of standard annuity rates specified by the Financial Services Authority (FSA); once the drawdown has commenced, the permitted maximum can be reviewed annually, and the income adjusted up or down if desired. (These rules allow a tax-free cash sum to be taken *without* an immediate taxable income: one simply chooses zero as the initial rate of income drawdown.) Pension income must be "secured" at age 75 at the latest, in one of two ways: either the residual fund can be used to purchase an annuity, or alternatively, an "Alternatively Secured Pension" (ASP) can be taken. ASP is essentially the same concept as income drawdown, but with more stringent limits on the income, which must be between 55% and 90% of standard annuity rates specified by the FSA. The maximum income can be reviewed annually while the ASP is in payment.

APPENDIX B

Simulations of Loss Harvesting

A model of loss harvesting

To give an indication of the potential benefits of loss harvesting (and also of tax aware management in general), it is convenient to consider a mechanical strategy where the investor makes no subjective decisions, and trades only to harvest losses. The strategy is as follows –

- The investor starts at time zero with a portfolio of 50 shares all priced at 100.
- The time horizon when the investor will liquidate the portfolio is 15 years.
- The benchmark portfolio is managed entirely passively.
- The loss harvesting portfolio is managed passively, except that once a month the investor reviews the portfolio and sells and repurchases any shares standing at a loss.
- We assume that the loss thereby generated can be set against other gains, arising either within the portfolio, from corporate events and rebalancing towards a benchmark index, or on the investor's other assets outside of the passively-managed, loss-harvested portfolio. The loss therefore generates a tax benefit.
- Our convention to account for the tax benefit of the loss (calculated as loss x tax rate) is to reinvest this amount immediately into the same share. (The time delay in payment of tax is ignored.)
- Every 6 months, dividends received in the past six months are invested into a new share.
- In both portfolios, corporate events and index rebalancing lead to one of the original 50 shares share being realised for cash and replaced every 8 months (with 50 shares, this represents a passive turnover of 3%pa).
- After 15 years, the two portfolios are liquidated, and tax paid on the gains on liquidation, with allowance for either indexation (for a corporate investor) or taper relief (for a personal investor).
- The benchmark and loss-harvested portfolios operate in identical economic environments, in particular in each simulation the *same* random numbers are used to generate the evolution of both portfolios. All results presented are based on 1,000 simulations.
- Initially, we ignore transaction costs.

In practice it is not possible to harvest a loss by selling and immediately repurchase the same share, because of the 30-day rule²⁰ whereby a sale is matched first against any repurchase within 30 days. For example, suppose a share was originally purchased for 100p and now stands at 95p. If we sell the share for 95p and then immediately (or within the next 30 days) repurchase it at 95½p, (say), the sale at 95p is matched against the purchase at 95½p to give a tax loss of ½p; the base cost for the share now held is the original 100p, and the loss from 100p to 95p has *not* been harvested. However, in principle we can get round this by investing the sale proceeds in a different share with similar risk and return characteristics. In practice it might be difficult to reliably identify such a share; but the purpose of our model is to give some indication of the potential benefit from loss harvesting, rather than to represent a real world strategy.

²⁰ In US papers, the equivalent provision is usually referred to as the 'wash sale rule'.

Models for share returns

We now need a model for the capital and dividend returns from shares generally, and a model for the cross-sectional variation in returns from different shares. Initially we follow the approach of previous US authors (Arnott, Berkin & Ye, 2001b; Berkin & Ye, 2003), who use the Capital Asset Pricing Model parameterised as follows:

$$r_i = (1 - \beta_i)r_F + \beta_i r_M + \varepsilon_i$$

where

r_i = monthly log return on share i

β_i = CAPM beta of share i ; Following the US authors, the betas are assumed normally distributed with a mean of 1, a standard deviation of 0.3, but capped at -1 and +3.

r_F = risk-free interest rate = 0.00506 per month (effective rate 6¼%pa)

r_M = expected total return on market portfolio of shares = 0.01 (effective rate 12.7%pa)

ε_i = residual error for monthly market log return

The risk-free interest rate of 6¼%pa was chosen to approximately correspond to the long-term median value in the Wilkie model (see below). The expected market total return was then set at a value which generated similar median accumulations in the portfolio without loss harvesting to those in the corresponding portfolio under the Wilkie model (see below).

The volatility of the market log return (that is, the standard deviation of ε_i) is set to 0.0433, that is 15%pa. The cross-sectional volatility of the log returns on individual shares around the market return is set to 0.09, that is 31%pa.

The CAPM gives total returns, but we need to separate out the dividend stream to allow for reinvestment say every six months, which establishes a higher base cost for a small part of the portfolio. I allow for dividends of 0.0028125 per month (3.375%pa) by deducting this amount from the CAPM total return on each share, and creating one new share with a base cost of 0.0028125 x 6 x current portfolio value every six months.

Table B.1 shows the portfolio relatives after 15 years, with and without loss harvesting, and before and after the tax paid on liquidation. A tax rate of 30% (appropriate to a corporate UK investor) is assumed. Indexation of base costs for capital gains is ignored in this table, since the CAPM approach does not produce a series for monthly RPI.

Table B.1. Portfolio relatives after 15 years of loss harvesting, CAPM returns

	Quartiles for portfolio relatives		
	Lower	Median	Upper
Before liquidation:	1.11	1.13	1.15
After liquidation:	1.09	1.10	1.11

Basis:

UK corporate tax rates (nil on dividends, 30% on capital gains)

No indexation or taper relief

Other parameters as per text.

Wilkie model for RPI and share returns

The CAPM is probably not a good model for long term share returns, and also does not provide the Retail Prices Index (RPI) figures which we need (at least in the case of a corporate investor) to include indexation allowance on base costs when calculating the capital gains on liquidation. A more suitable model for these purposes is an actuarial model known as the Wilkie model (see eg Wilkie, 1995). This gives separate series for RPI, share dividends and share yields (and hence share prices, derived as dividends/yields).

The Wilkie model produces only annual figures for each variable, but we need share prices at monthly intervals. These can be created by calculating the log share price at the end of each year (from dividends/yields), and then constructing Brownian bridges for monthly log share prices. The Brownian bridge from one year-end value to the next is constructed by generating a random walk 12 monthly steps forward from the starting annual figure, calculating the adjustment required to get from the end-point of this random walk to the required next annual figure, and then distributing this adjustment equally over the 12 monthly steps. (For a full description of this technique see Appendix D of Wilkie et al, 2003.)

The Wilkie model is parameterised as in Wilkie's 1995 paper, except that the long-term central value for the share dividend yield, YMU, is multiplied by 0.9 to reflect "actual" rather than "gross" dividends receivable since the removal of tax credits in 1999 (ie $YMU = 0.0375 \times 0.9 = 0.03375$). I experimented with changing the force of inflation, QMU, from 0.047 (the estimate based on long-term data in the 1995 paper) to 0.025, but it made little difference to the results. To start each simulation, I used 'neutral' initial conditions (Wilkie, 1995); the alternative would have been to use current market conditions, but since the Wilkie model says that shares are currently rather expensive, this approach might tend to exaggerate the benefits of loss harvesting. The Brownian bridge for monthly log share prices is based on a monthly standard deviation of 0.0433 (ie 15% pa). I also need to model the cross-sectional variation of monthly log price changes in individual shares around the log price change in the overall shares index. I assume that the cross-section is a normal distribution with a standard deviation is of 0.09 (ie 31%pa). This is in on top of the variability in the market index.²¹

Table B.2. Corporate investor: portfolio relatives after 15 years of loss harvesting, Wilkie returns

	Quartiles for portfolio relatives		
	Lower	Median	Upper
Before liquidation:	1.20	1.23	1.26
After liquidation:	1.14	1.17	1.20

Basis: as for Table B.1.

It can be seen that if the Wilkie model with Brownian bridges is used for monthly share price changes, the estimated gain from loss harvesting is somewhat higher than if the random walk model is used. For either model, the portfolio with loss harvesting is likely to be more than 10% ahead after 15 years.

²¹ These figures may seem on the high side, but my impression is that the portfolios of active investors tend to be concentrated in mid-cap and smaller company shares, where volatility does tend to be higher.

Further experiment showed that the results are not very sensitive to the assumed CAPM market rate of return, or Wilkie central inflation value QMU (which drives all the other variables in that model). The most critical parameter seems to be the cross-sectional standard deviation of individual log share price returns, that is in CAPM terms, the idiosyncratic risk.

Table B.3 shows the Wilkie model figures from Table B.2 adjusted for the statutory indexation of base costs in line with RPI when calculating capital gains (but not losses) on liquidation. The RPI figures used in the indexation are based on the Wilkie model, with monthly figures between each year end created by a Brownian bridge with standard deviation $1/\sqrt{12}$ times the Wilkie model parameter for annual standard deviation of inflation.²² As expected, indexation slightly reduces the benefit of loss harvesting; this is because the portfolio without loss harvesting has lower turnover, the shares held at liquidation are more aged and hence there is slightly more benefit in this portfolio from indexation than does the loss-harvested portfolio.

Table B.3. Corporate investor: portfolio relatives after 15 years of loss harvesting, Wilkie returns, with allowance for indexation of capital gains on liquidation

	Quartiles for portfolio relatives		
	Lower	Median	Upper
Before liquidation:	1.19	1.22	1.26
After liquidation:	1.13	1.16	1.18

Basis: UK corporate tax rates (nil on dividends, 30% on capital gains)

Table B.4 shows figures from the Wilkie model but with a primary tax rate of 30% and allowance for taper relief, based on the holding period for each share, on capital gains on liquidation. Taper relief has also been allowed on any capital losses on liquidation (but not on earlier losses generated from loss harvesting).²³ The portfolio relatives before liquidation are similar to those in tables 2 and 3. The reduction in portfolio relatives after liquidation is very slightly smaller than in Table B.3. This is probably because under indexation (Table B.3), the effective tax rate on liquidation of shares in the portfolio without loss harvesting continues to reduce the longer shares are held. Under taper relief (Table B.4), once a share has been held for 10 years and so qualified for maximum taper relief, any further holding period does not reduce the effective tax rate.

Table B.4. Personal investor: portfolio relatives after 15 years of loss harvesting, Wilkie returns, with allowance for taper relief of capital gains on liquidation; $t = 0.3$.

	Quartiles for portfolio relatives		
	Lower	Median	Upper
Before liquidation:	1.19	1.22	1.24
After liquidation:	1.14	1.16	1.18

²² Wilkie et al (2003) suggests using “Ornstein-Uhlenbeck bridge” for a quantity such as RPI where the annual series is AR(1), but for my present purposes any increased accuracy would be spurious.

²³ This is an approximation: the correct treatment is that losses on one share are set against gains on another share (or any other asset) chosen by the investor, and taper relief is applied to the net gain after offset of losses. Our assumption that taper relief applies to losses effectively assumes that the losses are set against gains on assets which had a similar holding period. In practice one would set losses against gains with the shortest holding period (equivalent to smallest amount of taper relief), and so our assumption may slightly under-estimate the benefits of loss harvesting.

Basis: as for Table B.3. (NB corporate tax rates used to aid comparison with earlier tables, even though this table shows taper relief.)

Note that the 30% rate of tax used in Table B.4 has been chosen to make the pre-liquidation figures comparable to those in earlier tables. The actual rate of tax for a personal investor is 40%. Table B.5 shows the same figures as in Table B.4, but with the correct tax rate of 40% for a personal investor. It can be seen that the portfolio relatives before liquidation (and hence also after liquidation) are higher than in earlier tables. This is because the higher rate of tax makes loss harvesting more advantageous.

Table B.5. Personal investor: portfolio relatives after 15 years of loss harvesting, Wilkie returns, with allowance for taper relief of capital gains on liquidation; $t = 0.4$.

	Quartiles for portfolio relatives		
	Lower	Median	Upper
Before liquidation:	1.31	1.36	1.42
After liquidation:	1.23	1.27	1.32

Basis:

- Personal tax rates (40% on tapered capital gains, 25% on dividends)
- Taper relief allowed on gains (and also on losses – see footnote).

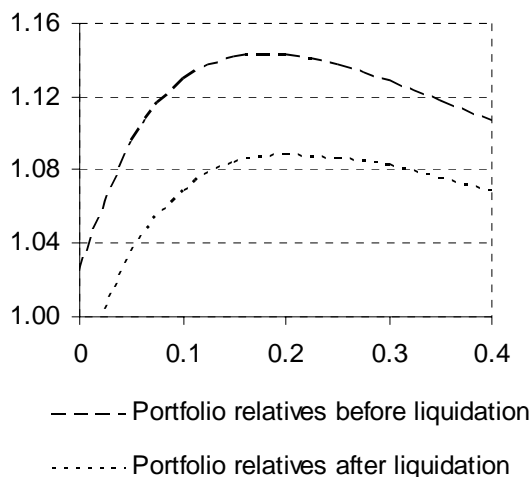
What about transaction costs?

The US authors' assessment of loss harvesting ignores transaction costs. This over-states the benefit of loss harvesting, and is particularly suspect for the UK because of stamp duty on purchases at 0.5%. In addition, in the UK commission on both purchases and sales might be 0.25% each, and bid-offer spreads perhaps 1%.²⁴ If we allow for these costs, it is no longer optimal to realise *any* loss; there is a trade-off between transaction costs and tax benefit obtained, and we should realise only losses above a certain size. Suppose we follow a rule to realise only losses greater than some threshold fraction f of base cost. We can then compare median terminal wealth relatives for different values of f .

It is convenient to present results as a graph of median terminal wealth relatives plotted against f . Figure B.1 is the plot for a corporate investor ($t = 0.3$ and allowance for indexation, same basis as Table A.3). Figure B.2 is the plot for a personal investor ($t = 0.4$ and allowance for taper relief, same basis as Table B.5). Comparing with Tables B.3 and B.5, the benefits of loss harvesting are (as expected) lower with transaction costs; but with suitable choice of f , they are still worth having. f just under 0.2 appears to give the best results for both corporate and personal investors. The optimal value of f tends to be slightly lower for a personal investor than a company (because the tax benefit on each loss is larger: $t = 0.4$ instead of $t = 0.3$), but the sensitivity of results to f in the range $0.1 < f < 0.3$ is fairly low.

²⁴ The bid-offer spread may seem a little low, but the extension of SETS to all of the FTSE All Share and much of the AIM index has reduced effective spreads: a patient dealer can now often buy at the bid and sell at the offer.

Figure B.1. Corporate investor: portfolio relatives as a function of threshold fraction f after 15 years of loss harvesting




Basis:

UK corporate tax rates (nil on dividends, 30% on capital gains, indexation of gains)

Transaction costs: 0.5% stamp duty, 0.25% commission either side, 1% bid-offer spread.

Figure B.2. Personal investor: portfolio relatives as a function of threshold fraction f after 15 years of loss harvesting



Basis:

UK personal tax rates (40% on tapered capital gains, 25% on dividends)

Taper relief allowed on gains (and also on losses – see footnote to Table A.5).

Transaction costs: 0.5% stamp duty, 0.25% commission either side, 1% bid-offer spread

It is interesting to note that without the benefit of simulation, one's first conjecture for a reasonable strategy might be "realise any loss where the expected tax benefit exceeds the round-trip transaction costs." This corresponds to $f \approx 0.06$ for a corporate investor, and $f \approx 0.04$ for a personal investor (lower for a personal investor because of the larger tax benefit on each loss: $t = 0.4$ instead of 0.3). Viewing each month in isolation, this strategy harvests the maximum tax benefit immediately available from that month; but it is not an optimal multi-period strategy. This seems to be because by selling a share at a small loss as soon as the loss emerges in month n , you forgo the chance to sell the *same* share (for roughly the same transaction cost) when the loss may have grown larger, in month $n+1$, $n+2$, $n+3$,...etc. It is better to wait until the loss is of more reasonable size. This is another example of the benefits of judicious lethargy.

The results in Figures B.1 and B.2 depend on our *ad hoc* model for cross-sectional share price returns. We have not estimated this model, and to do so would be a very substantial project in itself. However, as noted earlier (see comments on Table B.2), the cross-sectional standard deviation of individual log share price returns around the Wilkie "market" return seems likely to be one of the most sensitive parameters, and so a sensitivity test on this parameter seems sensible. I therefore reran the calculations for Figures B.1 and B.2 with monthly cross-sectional standard deviations of 0.04, 0.09 (the central case), and 0.14, which correspond roughly to 15%, 30% and 50% annualised figures. The optimal value of f was usually about 0.2, and was always in the range 0.15 to 0.25; this provides some comfort that using 0.2 would be reasonably robust. If transaction costs are lower, then as one would expect the optimal f is lower: for example, if bid-offer spreads are about 0.5% rather than 1%, the optimal f is close to 0.15 rather than 0.20.