THE APPLICATION OF DISCOUNT FUNCTION TECHNIQUES TO GOVERNMENT DEBT MARKETS:

With special consideration of tax and market segmentation problems

Thesis for PhD in Economics
University of London
September 1990

Candidate: **Dino Bibolini**
Supervisor: **Charles A. E. Goodhart**
College: **London School of Economics and Political Science**

ABSTRACT

Most empirical models of bond markets developed over the past quarter of a century have made use of discount function methodologies. Whilst representing a significant technical refinement, the use of discount functions is fraught with problems. These may, for convenience, be classified as relating either to "functional form" or to "market imperfections".

The choice of functional form usually emphasizes either a specified economic model or perceived regularities in empirical data. The formulation used in this study, while not too dissimilar from those that emerge from certain arbitrage models of the term structure, is designed with a view to being better behaved and easier to compute than others currently in use. Its wide-ranging applicability is illustrated by fitting the same model to two relatively simple cases - the German Bund and Dutch government bullet markets - and then extending it to the case of the UK gilt market.

Market imperfections cannot always be accommodated within a discount function model. This study focuses on two types of imperfections: segmentation or "clientele" effects, and taxation. Both these imperfections appear to be present in the gilt market, and the discount function methodology is adjusted so as to model this market by means of three representative investors endowed with different tax regimes and net of tax discount functions. The effective rates of tax and implied market segmentation resulting from the estimation appear to conform well to other empirical and anecdotal evidence. The main theoretical shortcoming of such an approach is the "buy-and-hold" assumption used to calculate tax payments. However, a partial test of the significance of this weakness concludes that it may be empirically unimportant in the case of gilts.

Finally, the framework thus developed is utilized to discuss various aspects of supply and demand influences on the gilt price structure.

TABLE OF CONTENTS

TABLE OF FIGURES

TABLE OF TABLES

I. DISCOUNT FUNCTIONS AND YIELD SURFACES

1.1 Introduction

The use of discount functions in empirical models of bond markets, which is nowadays widespread practice, is a relatively recent development, with the main body of existing research in this field dating back no more than 20 years. The principal reason for the emergence of this methodology is the fact that the term structure of interest rates or yields is not in general directly observable. The relevant data cover prices[1] of individual assets (deposits, bills, bonds, etc.) which may - or, indeed, may not - reflect the existence of a term structure, and usually incorporate complex unobservable attributes such as credit risk.

Since the term structure is unobservable it can at most be estimated. Therefore, the numerous tests and applications of economic theory involving results of such an estimation are dependent upon the approach followed. This is important because different assumptions can produce very different results, as is apparent from the discussion of the UK gilt market in chapter III.

The first two chapters of this study focus on the use of discount function methodologies for modelling bond markets when there are no tax effects. After a brief survey of the relevant literature, contained in the next section, a number of theoretical results are presented and discussed. The theory is followed up in chapter II by empirical evidence based on case studies of the German Bund and the Dutch government bullet markets. These bond markets are well-suited to the

1. Or, equivalently, nominal rates of interest, yields, etc.

purpose. Although the former was subject briefly to withholding tax after 1 January 1989, it was dominated during the period covered in the model estimation by investors who were in the main concerned about gross returns.[1] The Dutch government bullet market was and still is an example of a market with virtually no tax effects.

The following five chapters deal with the much more complex case of the UK gilt market, which is characterized by significant tax effects and segmentation. After describing how the basic discount function approach might be extended to cope with this situation, the results of empirical work are presented and critically evaluated. In the process, a number of tangential topics of significant interest are examined, such as the so-called FOTRA effect, or the shocks arising from the announcement of new issues.

1. Even after the new tax was announced, the relative pricing of bonds was, for a number of reasons, only slightly affected by the tax.

1.2 Discount function literature

Term structure theory dates back many decades, but serious attempts at estimating discount functions before the 1970s were few and far between. This is at least in part due to the fact that the more primitive versions of term structure theory were ill-suited to cope with the subtleties of real fixed interest markets. Most models of the term structure have been couched in terms of a pricing rule for tax exempt discount bonds, implicitly assuming the existence of a complete set of such bonds, and paying little attention to the general relationship between term structure and bond prices.

Bond market models incorporate varying degrees of structure, ranging from the purely descriptive type to the asset pricing variety based on a detailed general equilibrium framework. The former extreme consists of nothing more than a contrived mathematical function fitted in a price or yield space;[1] the results obtained from the application of such a method might be inconsistent with any sensible economic theory. Conversely, proper economic models, while potentially providing a more accurate description, incorporate prior information that might heavily bias the results and any subsequent use thereof.[2]

There are a number of specific factors that are generally believed to influence the price of a bond, some of which are the length of the loan, the "coupon", the "liquidity" of the issue, the credit worthiness of the issuer, and the taxes applicable to the deal. The use of discount functions, which would be straightforward in the case of a perfect bond

---

1. Or in any other one-to-one transform of the price, e.g. its reciprocal (usually called the "capital").

2. A trivial example of this would be a model of the UK gilt market which incorporated a risk premium positively related to volatility, and concluded that in 1987-88 either market participants' expectations of future interest rates were unreasonable or that a large proportion of bonds was heavily mispriced (see the chapters on the gilt market for further details).

market, is fraught with problems in practical situations. The existence of taxation is probably the single greatest nuisance, since the tax effects depend on circumstances such as the type of bond, the type of holder, or the prospective date of disposal, each of which can give rise to awkward difficulties. Another problem, which may be enhanced by differences in tax status between groups of investors, is market segmentation, i.e. the fact that certain bonds may be held exclusively by certain categories of investors. These issues are summarily discussed in the following subsections.

1.2.1 No taxation and no segmentation

These are the assumptions made in the more primitive models of bond markets. The resulting relationship between term structure and bond yields is discussed, inter alios, in Buse [7], Caks [8], Carr, Halpern and McCallum [11], Khang [30], Livingston and Caks [38], and Schaefer [48]. These studies assume complete bond markets[1] or an extreme form of expectations hypothesis, thus avoiding the fact that in real world situations the very existence of a discount function may be brought into question. The existence theorem outlined in section 1.3.2 is a simple extension of an analogous result described in Schaefer [47].

Early versions of term structure theory were closely linked to a form of expectations hypothesis. This appears in the literature in a number of guises, the most respectable of which are discussed in Cox, Ingersoll and Ross [18].[2] At a later stage the emphasis moved to modelling the "risk" or "liquidity" premium, culminating in a number of sophisticated continuous time models, such as those detailed in Brennan and Schwartz [3] or in Cox, Ingersoll and Ross [19]. Only in recent years have models of the term structure appeared that make allowance for taxation and segmentation.

1. A bond market is complete if it is possible to construct, by means of a portfolio of bonds, any cash flow contained in a bond belonging to that market.

2. See also Campbell [9] for an alternative view.

1.2.2 Segmentation

Although informal bond market segmentation theories have been around for many years, segmentation has become a major feature of the theory thanks mainly to the efforts of Schaefer [e.g. 45, 46, 47] at the turn of the 70s. This followed close on the heels of extensive research into the existence of segmentation or "clienteles" in asset markets generally and in equity markets in particular. Fixed income markets, however, thanks to the near-singularity of their return covariance matrix, were ideal candidates for identifying the existence of segmentation. Schaefer's approach is based on assuming that negative holdings are not allowed, thus ensuring the existence of an equilibrium in the case of investors with different tax rates.[1] A discussion of how constraints on investors' positions may be replaced with transaction costs and tax asymmetries is contained in Dermody and Prisman [23].

Whilst particularly useful in the case of differential taxation, such constraints could also be assumed in the general case. This is briefly explored in section 1.3.2, but did not appear empirically to be applicable to the case studies detailed in chapter II. However, segmentation without differential taxation has received little attention in the empirical economic literature, and is a topic of research that might deserve pursuing further.

1. See section 3.5.3 for more details.

1.2.3 Taxation

In most bond markets tax considerations are so important that they cannot be neglected without leading to absurd results. And yet, by and large, term structure theory was developed with scant consideration for the existence of taxes. This is probably due to the complexity of term structure models with taxation. Furthermore, since tax regulations are specific to a bond market and to a moment in time, it is impossible to produce a general model with taxation.

Although a general model may not be produced, there are obviously certain principles that are widely applicable. The effects of taxation usually depend on the different tax treatment of "principal payments", "income" and "capital gains". A discussion of the problems that arise from this diversity of cash flows is contained in Litzenberger and Rolfo [31]. Empirical work based on a model that assumes one representative investor endowed with some kind of "effective tax rates" is contained in Cramer and Hawk [21], Jordan [29], Livingston [34, 36, 37], McCulloch [41], and Robichek and Niebuhr [43]. Different tax brackets within a segmented market are modelled in the previously mentioned studies by Schaefer, in Hodges and Schaefer [28], and Ronn [44].

A major problem with allowing for taxation is that, even if it is assumed that the future tax regulations are known in advance, the tax payments attached to a bond usually depend on (maybe successive) holders' trading strategies and on future states of the world. For example, as explained in Constantinides [14, 15], Constantinides and Ingersoll [16], and Constantinides and Scholes [17], the buy-and-hold or buy-and-sell strategies usually used to derive net of tax pricing models are sub-optimal if capital gains are taxed or capital losses deducted only when realized. A discussion and empirical investigation into these issues in a spirit similar to that of Litzenberger and Rolfo [32] is contained in chapter V.

The analysis of the gilt market of chapters III and IV is closest in approach to the work of Schaefer, although it differs from this in certain respects, such as the *joint* estimation of discount functions for different tax brackets, and the functional form chosen. Chapter IV also contains an analysis of how tax rates on gilts implied by the model changed in response to expected and actual tax changes in the UK 1988 budget. There is little empirical literature at present illustrating this kind of effect.

1.2.4 Functional form

Finding functional forms suitable for modelling the term structure is particularly difficult. This is reflected in the literature by the fact that, with the exception of spline methods, which are extensively used in the more recent studies, virtually every author adopts a different approach. A thorough discussion of the use of splines for term structure estimation is contained in Shea [49, 50].[1]

The approach used in this study is somewhat in the spirit of Nelson and Siegel [42], in that the functional form, which is not a spline, is the solution to a linear differential equation. As well as providing a flexible and smooth curve, this has the technical advantage of being easy to implement and fast to compute. The model differs from any other currently in the literature in that it incorporates a Bayesian prior, which helps to ensure that the discount function is well-behaved.

1. A similar approach is discussed in Vasicek and Fong [52].

1.3 Basic definitions and assumptions

1.3.1 Bond

A bond is defined as a finite sequence of cash payments, the amount and timing of which is known in advance.[1] In accordance with common practice it is assumed that the timing of a payment is adequately described by the date on which it is due. The range of dates on which a payment is possible is assumed to be limited to a (large) finite set. Under these conditions a bond is, from a mathematical point of view, a vector in the positive orthant of a finite-dimensional space:

$$b \in \Omega = \Re_+^T$$

where b denotes the bond, Ω the bond space, T the date space, and $\Re_+$ the set of non-negative reals.

1. This excludes instruments that pay a variable amount of interest, or with unknown redemption date. Furthermore, it rules out credit risk: the models developed in this study are targeted at homogeneous and high-grade government bond markets.

1.3.2 Bond market

In this study a bond market is modelled as a kind of Arrow-Debreu market, where the commodity space is the bond space Ω. The bond market discount function is the vector of (shadow) present values of one monetary unit payable on each possible payment date. The discount function is therefore a price vector for Ω which is normalized so that the unit price of cash payable currently is one. If v denotes the discount function, the (theoretical) price p of a bond is, in standard notation:

$$p = b'v.$$

Even in the absence of taxation, the assumption that a bond market may be modelled in this way is very strong due to the existence in practice of circumstances capable of causing market segmentation. The following set of formal results illustrates the point.

Let $B = \{b_1, b_2, \ldots b_H\}$ be the array of bonds in a given market and $\alpha_0 \in \Re_+^H$ be the vector of total amounts in existence. Consider the following assumptions:

A1. B and α_0 are exogenously determined.[1]

A2. A competitive equilibrium exists for the relevant economy such that each i-th investor chooses a portfolio $\alpha_i \in \Re^H$ and a payment for the portfolio of $m_i \in \Re^1$, given a vector of bond prices $p \in \Re^H$.

---

1. I.e. there is a sole issuer (e.g. the government) whose behaviour can be taken to be exogenous to the model. Individuals have no ability to modify B or α_0 by issuing new bonds or "selling short" existing ones.

A3. There exists an i-th investor who's investment choice can be described by an ordering $\succ_i$ defined on the set $X=\{x \in \Re^T \mid \exists \alpha \in A, \exists m \in \Re^1 : x = B\alpha - e_1 m\}$, where $A = \{\alpha \in \Re^H \mid 0 \le \alpha \le \alpha_0\}$ and e_1 is the first T-dimensional unit vector,[1] and satisfying the strict monotonicity property:

$$x_1, x_2 \in X, x_1 > x_2 \Rightarrow x_1 \succ_i x_2.$$

A4. There exists an i-th investor who's investment choice can be described by an ordering $\succ_i$ defined on X as in A3, satisfying the time-preference property:[2]

$$x_1, x_2 \in X, \varepsilon \in \Re^1_+, x_1 > x_2 - (e_{t_1} - e_{t_2})(t_1 - t_2)\,\varepsilon \Rightarrow x_1 \succ_i x_2.$$

A5. The α_i chosen by the i-th investor according to assumption A3 or A4 satisfies $\alpha_i \ll \alpha_0$ (segmentation condition).

A6. The α_i chosen by the i-th investor according to assumption A3 or A4 belongs to the interior of A (non-segmentation condition).[3]

---

1. The notational convention adopted for vector inequalities is:

$(y \ge x) \Leftrightarrow (\forall i \in I : y_i \ge x_i)$

$(y > x) \Leftrightarrow (\forall i \in I : y_i \ge x_i, \exists i \in I : y_i > x_i)$

$(y \gg x) \Leftrightarrow (\forall i \in I : y_i > x_i)$

where $y, x \in \Re^I$.

2. This assumption, which implies strict monotonicity, is not much stronger, since it only requires in addition that money can be stored at no cost.

3. This assumption clearly implies the previous one.

Theorem 1 (existence of a discount function)

Under A1-A3 and A6 there exists a discount function, i.e. a vector $v \in \Re_+^T$ such that $p = B'v$ and $v_1 = 1$. Under A1-A3 and A5 there exists a segmented discount function, i.e. a vector $v \in \Re_+^T$ such that $p \geq B'v$ and $v_1 = 1$.

Proof

Consider the set $Y_i = \{y \in \Re^T \mid \exists\alpha \in A, \exists m \in \Re^1 : m = p'\alpha,$ $y = B(\alpha - \alpha_i) - e_1(m - m_i)\}$. Clearly Y_i is a convex and closed set. Furthermore $Y_i \cap \Re_+^T = \{0\}$, since otherwise the i-th investor could choose α, m with $B\alpha - e_1 m > B\alpha_i - e_1 m_i$, and therefore $B\alpha - e_1 m \succ_i B\alpha_i - e_1 m_i$, which is impossible given the assumed choice of α_i, m_i. So there must exist a hyperplane separating[1] Y_i and $\Re_+^T - \{0\}$, and supporting Y_i through the origin:

$$\exists v_i \in \Re^T - \{0\}, \forall y \in Y_i, \forall e \in \Re_+^T - \{0\} : y'v_i \leq 0, e'v_i > 0.$$

Since then $v_i \gg 0$, one may take $v_{1,i} = 1$ with no loss of generality. Furthermore:

$$\forall\alpha \in A, \exists m \in \Re^1 : m = p'\alpha, [B(\alpha - \alpha_i) - e_1(m - m_i)]'v_i \leq 0$$

and therefore:

$$\forall\alpha \in A : (B'v_i - p)'(\alpha - \alpha_i) \leq 0.$$

The proof is completed by comparing this expression with A5 and A6. ■

1. By a variant of the separating hyperplane theorem strict separation is obtainable for $\Re_+^T - \{0\}$.

Theorem 1 is central to discount function model building. The non-segmentation condition is unlikely in practice to fail because the whole of a bond in existence is held by only one investor, but could easily fail because there isn't an investor that holds a strictly positive amount of all bonds. In this case it might be possible to partition the market into segments such that a segmented discount function applies to each one. If the segmented discount functions are denoted by v_j, then $p_h \geq b_h' v_j$, where strict inequality can occur only for the bonds not belonging to the j-th segment.

Theorem 2 (monotonicity of the discount function)

If A3 is replaced by A4 in the previous theorem, the (segmented or non-segmented) discount function can be chosen to be monotonically non-increasing.

Proof (outline)

Let $S = \left\{ s \in \Re^T \mid \exists \varepsilon \in \Re^1_+, \exists t_1, t_2 \in T : s > -(e_{t_1} - e_{t_2})(t_1 - t_2)\,\varepsilon \right\}$.

The separation then applies to Y_i and S, and the proof is otherwise virtually identical to that of the previous theorem, except for the additional result that, since:

$$\forall \varepsilon \in \Re^1_+, \forall t_1, t_2 \in T : -[(e_{t_1} - e_{t_2})(t_1 - t_2)\,\varepsilon]' v_i \geq 0$$

one gets:

$$(v_{t_1,i} - v_{t_2,i})(t_1 - t_2) \leq 0. \;\blacksquare$$

Theorem 2 provides a very weak characterization of a discount function. In view of the following Theorem 3, much stronger constraints are usually required for actual estimation.

Theorem 3 (uniqueness of the discount function)

If a discount function v exists for a price vector $p \neq 0$, it is unique if and only if B is of full row-rank, i.e. if the bond market is complete.

Proof

If B were not of full row-rank there would exist a non-null vector k such that $B'k = 0$. Then $\exists \varepsilon \in \Re^1 : \bar{v} = (v + k\varepsilon)/(v_1 + k_1\varepsilon) \neq v$ and $\bar{v}$ would be another discount function. Conversely, if B were of full row-rank while v_1 and v_2 were two different discount functions, then $B'(v_1 - v_2) = 0$, a contradiction. ■

Theorem 3 is of little comfort in practice since the uniqueness conditions are unlikely to be satisfied. If preferences were continuous, convex and sufficiently smooth,[1] one could choose v to be the vector generating the hyperplane separating Y_i and $S_i = \{y \in \Re^T \mid \exists \alpha \in A, \exists m \in \Re^1 : y = B(\alpha - \alpha_i) - e_1(m - m_i), B\alpha - e_1 m \succ_i B\alpha_i - e_1 m_i\}$. Since S_i and Y_i in general vary from investor to investor, unless a representative investor is assumed, this approach doesn't provide a unique discount function. However, even a representative investor isn't sufficient to remedy the lack of statistical identification of the discount function model. As discussed in section 1.4 some additional information is therefore required.

1. I.e. generated differentiable indifference curves. One may wish to restrict this argument to a neighbourhood of $B\alpha_i - e_1 m_i$.

1.3.3 Yield

The traditional method of assessing the value of a bond is based on the "yield". This is the internal rate of return on the money invested in the bond. In standard notation, if p is the market price of the bond, b_n the n-th cash flow and t_n the time (in years) to the n-th cash flow, the (annual) yield[1] is defined as the solution y to the following equation:

$$p = q(y, b) = \sum_n b_n (1 + y/100)^{-t_n}.$$

The shortcoming of the yield as a general criterion for bond valuation was pointed out long ago: all cash flows are valued on the basis of the same discount rate, and this is inconsistent with a variable term structure. The reason why this approach is still so commonly used by investors and dealers alike is that a yield can be *calculated* from an individual price while a set of different discount rates for each cash flow usually must be *estimated.* Yields are used extensively in this study despite the focus on discount functions because they are often more easy to understand than prices and provide a more straightforward comparison with certain other models.

1. Hereafter the notion of yield used will, unless expressly stated, be that of f-annual yield, obtained by applying to the annual yield the transformation $100 f\,[(1 + y/100)^{1/f} - 1]$, where f is the coupon frequency.

1.3.4 Yield surface

Postulating a discount function model has strong implications for the structure of yields. To illustrate this fact it is first useful to make a simplification: given the coupon frequency,[1] a bond can usually be described to a good approximation by its coupon and life. So, in mathematical notation, taking frequency as a known parameter and denoting coupon and life by c and t, one may write:

$$b = b(c, t).$$

Since, as shown previously, the yield is by definition a function of bond and price,[2] one may also write:

$$y = y(p, b).$$

Finally, if the price of a bond is determined in accordance with the (non-segmented) discount function approach, it is a function of bond and discount function, i.e.:

$$p = p(v, b).$$

Therefore, substituting appropriately, one gets:

$$y = y(p(v, b(c, t)), b(c, t)) = z(v, c, t)$$

i.e. the yield depends on discount function, coupon and life.

---

1. The coupon is the amount of interest paid on the bond in the course of one year. The frequency is the number of instalments in which the coupon is paid.

2. It is shown that, under realistic assumptions, the equation defining the yield always has a unique solution.

For any given v^* the function $y = z(v^*, c, t)$ describes what is known as the "yield surface". It is easily shown that a yield surface uniquely defines a discount function. However, not all functions can be yield surfaces according to this approach; in fact, roughly speaking, given one slice of the surface, e.g. a "coupon yield curve" obtained by fixing the coupon at an arbitrary level c^* to get the function $y = z(v^*, c^*, t)$, the rest of the surface is also determined. This set-up has, therefore, strong implications for the structure of yields.

A graph of a yield surface that has been calculated by taking an estimated par yield curve,[1] working out the implied discount function and from that obtaining the rest of the surface is given below. The par yield curve generating the graph is also drawn as a dark line across the surface. This corresponds to the intersection of the surface with the plane where yield is equal to coupon.[2]

---

1. The par yield curve corresponding to the yield surface $y = z(v^*, c, t)$ is the function:

$y = z(v^*, c(v^*, t), t)$

where $c = c(v^*, t)$ is the solution to the equation:

$100 = p(v^*, b(c, t))$

i.e. describes the path over the yield surface where the price of the bond is equal to par.

2. In this stylized setting, if a bond is priced at par on a coupon payment date, the yield is equal to the coupon. In between coupon payment dates, this equality is only approximately true.

Yield surface

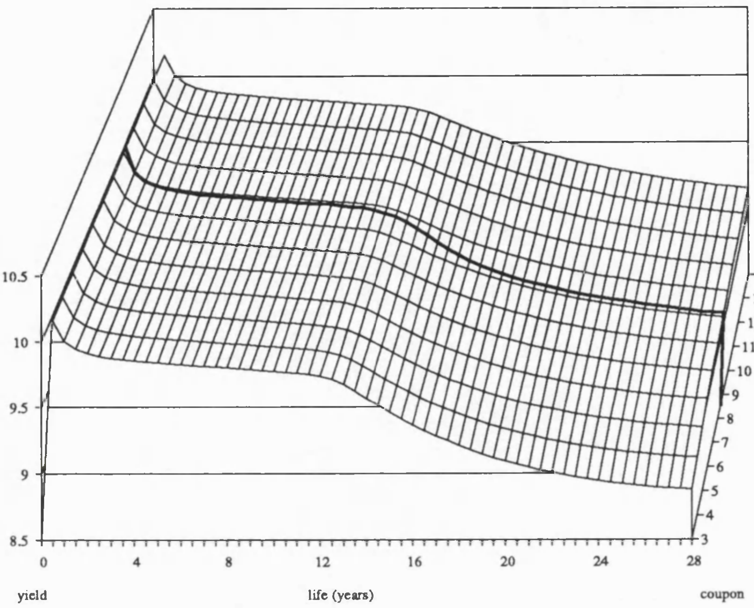

The graph above provides some insight into a feature of empirical yields commonly known as the "duration effect".[1] This consists in the tendency for high-coupon bonds to yield more than low-coupon bonds in what may be loosely described as a "downward-sloping yield curve environment", and vice versa in an "upward-sloping yield curve environment". The reason for this is related to the notion of duration, from which the name derives.

---

1. The duration effect is a special case of "coupon effect". The coupon effect is the tendency, common to most bond markets, for yields to depend on coupon as well as on maturity. The dependency may be ascribed to the duration effect, or may be due to other factors, usually taxation. An example of well-defined "taxation effect" is discussed in the analysis of the UK gilt market.

1.3.5 Duration effect

Duration (Macaulay duration) is defined as the weighted average life of the bond's cash flows, where the weights are the present values of the payments calculated using the bond's yield as discount rate. An equivalent and often more useful definition is that duration is the point percentage change in price corresponding to a unit change in the instantaneous yield $\bar{y}$:[1]

$$d(y, b) = -100\, \partial \ln q(y, b)/\partial \bar{y}.$$

A commonly held belief on the part of market participants is that bonds should tend to have similar yields when they have similar durations rather than lifes. More precisely:

$$y(b) = f(d(y(b), b))$$

where f(.) is a given function. This could help to explain the duration effect, but is inconsistent with the assumed discount function model. The latter implies that the yield depends on the profile of cash flows, but not that it is only a function of duration. It can actually be shown that the two conditions could be satisfied only if the term structure of yields were flat.[2] The reason for the belief lies mainly in the portfolio immunization techniques used by many fund managers, and particularly in the simplistic assumptions underlying those techniques.[3]

1. The instantaneous yield is obtained by applying the transformation $100 \ln(1 + y/100)$ to the annual yield, i.e. is analogous to the instantaneous rate of interest.

2. For a proof of this result see the appendix.

3. This point is further discussed in the appendix.

It can be shown that, in a discount function model:[1]

i) the sign of the duration effect at a given maturity is determined independently of coupon level

ii) the duration effect is positive at a given maturity if the yield, for any coupon level, is lower than the annuity yield[2] at the preceding maturity. The duration effect is negative if the converse is true.

1. For a proof of these results see the appendix.

2. The annuity yield is the yield on a bond with a principal repayment of zero. It is shown to be equal to the limit of the yield of a coupon bond when the coupon tends to infinity.

1.4 A simple discount function model

The basic discount function model used in this chapter and the next is of the form:

$$\begin{cases} p_h = b'_h v + \varepsilon_h \\ E\varepsilon_h = 0, E\varepsilon_h \varepsilon_k = \begin{cases} \sigma_h^2 & \text{if } h = k \\ 0 & \text{otherwise} \end{cases} \end{cases}$$

i.e. is a classical linear model. The main difficulty with this model is that it is over-parameterized, since in most realistic situations the dimension of v, i.e. the number of dates on which a cash payment can be made, is much greater than the number of bonds that can be used for the estimation.[1]

One way around the problem is to restrict the parameter space by re-parameterizing the discount function, i.e. to postulate:

$$v = v(\beta)$$

and then estimate the (possibly non-linear) model equation:

$$p_h = b'_h v(\beta) + \varepsilon_h.$$

1. The US Treasury market comes close to having enough bonds relative to the number of payment dates.

Another method, which effectively amounts to enlarging the observation space, consists in adding a Bayesian prior to the original equation,[1] to obtain:[2]

$$\begin{cases} p_h = b_h' v + \varepsilon_h \\ f_l(v) = \mu_l \\ E\varepsilon_h, \mu_l = 0, E\varepsilon_h \mu_l = 0, E\varepsilon_h \varepsilon_k, \mu_l \mu_m = \begin{cases} \sigma_h^2, \hat{\sigma}_l^2 & \text{if } h = k, l = m \\ 0 & \text{otherwise.} \end{cases} \end{cases}$$

On empirical grounds, the best solution appears almost invariably to be a mixture of the two approaches. The model used is therefore of the form:

$$\begin{cases} p_h = b_h' v(\beta) + \varepsilon_h \\ f_l(v(\beta)) = \mu_l \\ E\varepsilon_h, \mu_l = 0, E\varepsilon_h \mu_l = 0, E\varepsilon_h \varepsilon_k, \mu_l \mu_m = \begin{cases} \sigma_h^2, \hat{\sigma}_l^2 & \text{if } h = k, l = m \\ 0 & \text{otherwise.} \end{cases} \end{cases}$$

The choice of functions v(.) and f_l(.) should take into account the fact that, under the assumptions stated in section 1.3.2:

i) a discount function can be chosen to be monotonically non-increasing

ii) $v_0 = 1$.

---

1. This is an extension of the mixed estimation methodology advocated by Theil and Goldberger (see e.g. Theil, Principles of Econometrics, 1971).

2. Care must be taken in choosing the prior to avoid inconsistencies with the discount function parameterization.

The choice of v(.) should also, possibly, be linear in β, to make for computational simplicity. A method that, on the basis of empirical trials, seems to give consistently satisfactory results is to use an expression of the form:

$$v_t(\beta) = \sum_{k=1}^{K} \beta_k \exp(-r_k t) + \left(1 - \sum_{k=1}^{K} \beta_k\right) \exp(-r_{K+1} t)$$

where the rates r_k follow some pattern. For this purpose, a geometric progression is found to be convenient and adequate, although other patterns perform well. Since the results are almost invariably little affected by quite large changes in the set of rates used, the method adopted is to choose a starting value r_1 and a ratio r_0 on the basis of a large sub-sample of data, and then keep these fixed for the whole sample. This formulation appears to be capable of producing all sorts of humped shapes[1] and to be generally well-behaved at the extremes of the sample. No attempt has been made to provide a basis for this functional form in terms of a specified economic model of the term structure, although the fact that it is the solution to a linear differential equation suggests that this might be possible.

The Bayesian prior functions $f_l(.)$ are always taken to be of the form:

$$f_l(v_t) = (\Delta \ln)^2 v_{t_l}$$

i.e. the prior probability distribution of forward instantaneous yields is chosen to be a geometric martingale.[2] The prior has generally not been found to be necessary to produce sensible results, and is never taken to be so strong as to make an appreciable difference to the actual fit. Its main purpose is that, in the case of more complex situations,

1. E.g. with $k = 4$ it is usually easy to pick up 2 humps.

2. As for any prior, this choice is largely arbitrary. It is worth noting that it implies positive forward yields.

and in particular in the case of the gilt model, it speeds up convergence significantly by helping the iterations to find the right path to the optimum.

Another problem with this model is the estimation of the error variance. The simplest approach is to assume homoscedasticity. In practice this is likely to be inaccurate, since one would expect the pricing accuracy to be very much related to the bid-offer spread. The latter is, in fact, the cost to the arbitrageur of one leg of the "round-trip" into and back out of a relatively cheap bond from a relatively dear one. The cost of the other leg depends on the bid-offer spread of the other bond(s) involved in the "switch", and therefore cannot be determined a priori. The bid-offer spread on a given issue depends on a number of factors, but an important one is volatility,[1] which usually increases with duration. The latter or some related variable may, therefore, provide a rough guide for the pricing error standard deviation. If the standard error were homogeneous in duration, using this result would be very similar to fitting the model in the yield space under homoscedasticity.[2]

---

1. I.e. the annualized standard deviation of the return on an investment of one unit.

2. Let ε_h^p and ε_h^y be the errors in price and yield respectively on the h-th bond. Then, to a linear approximation:

$$\varepsilon_h^p = [\partial p_h/\partial y_h]\varepsilon_h^y$$

where the derivative is computed at the current price. Therefore OLS in the yield space consists in minimizing:

$$\sum_h(\varepsilon_h^y)^2 \approx \sum_h(\varepsilon_h^p/[\partial p_h/\partial y_h])^2$$

which is approximately the same as WLS in the price space if it is assumed that:

$$\sigma_h^p = [\partial p_h/\partial y_h]\sigma^y.$$

If duration d_h were used instead of $[\partial p_h/\partial y_h]$, the result would be similar, since:

$$d_h = -[\partial p_h/\partial y_h]\,[100/p_h]\,[\partial y_h/\partial \bar{y}_h]$$

and $[100/p_h]\,[\partial y_h/\partial \bar{y}_h]$ is approximately constant.

In practice, there is little evidence that such an approximation is much more accurate than the simpler homoscedastic assumption. This point is discussed in greater detail with reference to the UK gilt market in section 4.5.

II. TWO CASE STUDIES

2.1 The German Bund market

2.1.1 Introduction

The West German bond market is the third largest in the world, but only a relatively small proportion of it consisting essentially of bonds (and notes) issued by the Federal government is actively traded.[1] There are three categories of government issues that are truly liquid by international standards; these are the Bundesanleihen, the Bundesobligationen, and the Bundeskassenobligationen. With the exception of one bond, which is neglected hereafter, all three categories consist of straight bullets.[2] But there are significant differences that cause these instruments to trade on separate markets.

Bundeskassenobligationen are notes issued with maturities of up to just under 6 years. Although actively traded over the counter, especially when recently issued, they are not officially quoted on the stock exchanges. Therefore the market is less transparent than in the case of Bundesanleihen or Bundesobligationen.

1. Much emphasis is put hereafter on the need for liquidity. This is for two reasons. Firstly, only if a bond market is liquid is it possible to obtain a data set that accurately reflects the prices at which market participants would have been prepared to trade at a specified moment in time. Secondly, a discount function model is ill suited to coping with differences in liquidity amongst bonds.

2. A "bullet" is the standard bond, consisting exclusively of coupon payments at fixed dates throughout its life and a principal repayment on a fixed redemption date. So, for example, a "callable" bond, i.e. a bond that may be redeemed early in certain circumstances, is not a bullet. A "straight" (bond) consists exclusively of coupons and principal repayment of fixed predetermined amounts. So, for example, a "floating rate note", i.e. a bond that pays variable amounts of interest on coupon payment dates, is not a straight.

Furthermore, Bundeskassenobligationen do not benefit after issue from the regulatory intervention of the Bundesbank (see below), which helps to maintain liquidity in all issues. An important consequence is that foreign investor participation is relatively scarce, whereas in recent years it has been a major determinant of the prices of Bundesanleihen. So Bundeskassenobligationen usually trade at a yield spread above comparable Bundesanleihen.

Bundesobligationen are notes issued with maturities of around 5 years on a tap basis. Their main distinctive feature was until recently the fact that foreign investors were not allowed to purchase them.[1] This limitation affects most of the sample period used for the model estimation.

Bundesanleihen, known more commonly as "Bund", are bonds issued usually with maturities of around 10 years. There are approximately 80 Bund in existence, with maturities of up to almost 30 years. The two longest maturity Bund are due for redemption in the year 2016 and are the only really long Bund. This leaves a large gap in the maturity spectrum between 10 and 30 years. The size of Bund issues is usually of DM 1-5bn, with the most recent ones at the higher end of the range. Recent issues are therefore relatively large by international standards, and this enhances their liquidity. In terms of overall size the Bund market is not, however, all that large, amounting to only about DM 150bn.[2]

1. This limitation has been suppressed with effect from 3 October 1988.

2. About half the size of the UK gilt market.

2.1.2 The market place

The Bundesanleihen market is carefully regulated by the Bundesbank, which intervenes daily on the stock exchanges to maintain liquidity and prevent erratic price movements. As a result, even issues that were made long ago and are therefore infrequently traded are priced in line with the rest of the market.

The important role played by the German central bank on the stock exchanges in determining the price of Bund should not, however, lead one to think that the nature of the market is so special as to make a comparison with the other markets covered in this study relatively meaningless. While the stock exchanges represent a point of reference and ensure liquidity for the less liquid Bund, the vast majority of the turnover is carried out over-the-counter, and a very sizeable proportion of this business is in London.[1]

The method of issue used for Bund deserves special mention since it affects the structure of yields in the 8-10 year maturity range. Bund are issued to a syndicate of banks which, in exchange for a re-allowance of 1.375%, underwrite the issue and undertake to place it with firm investors. If at any time up to one year after issue the Bundesbank purchases the bond below par in the course of its open market operations on the stock exchanges, it is deemed that the bond in question was badly placed, and the syndicate member responsible is obliged to repay to the Bundesbank most of the re-allowance (1%).

1. This is obviously the case for UK gilts, and is also the case for Dutch government bullets.

During the first year after issue Bund can be transmitted in two forms: free or restricted. When Bund are sold in restricted form the purchaser undertakes, in exchange for a lower price, to ensure that the bonds are not sold on to the Bundesbank below par within a year from issue. If a purchaser of Bund in restricted form doesn't fulfil this undertaking the seller can reclaim the discount given and will in turn himself be liable for any discount obtained. When Bund are sold in free form the purchaser assumes no such obligation.

Immediately after issue "free form" stock is relatively scarce, and trades at a lower yield. The relative scarcity of free Bund declines steadily, until after 1 year all Bund can be sold freely. Stock sold on the stock exchanges must always be in free form. This helps to ensure that, on stock exchange data, yields between 8 and 10 years form a hump, as holders of restricted Bund try to avoid selling, whereas holders of Bund that have recently become free often want to sell.

Another major factor contributing to the hump in the 8 to 10 year maturities is liquidity, which tends to depress yields of newly issued Bund, and also causes the hump to become particularly pronounced during rallies and in periods of intense speculative interest on the part of foreign investors. Finally, the hump is also enhanced by the fact that yields in the 9-10 year area are depressed by the big gap in the spectrum of maturities between 10 and 30 years.

2.1.3 Empirical results

The Bund model was estimated for the period 1987-88 using daily data obtained from Bloomberg (which closely reflect stock exchange fixings) and Reuters (which are stock exchange fixings). Up to 1 June 1988 only 15-20 Bund were used and the main source was Reuters, whereas thereafter the whole market (i.e. around 80 issues) was used and the main source was Bloomberg.

The discount function formulation that seems best to fit the data over the two year period is:

$$v_t(\beta) = \sum_{k=1}^{K} \beta_k/(1+r_k)^t + \left(1 - \sum_{k=1}^{K} \beta_k\right)/(1+r_{K+1})^t$$

where:

$K = 4$
$r_1 = 1\%$
$r_2 = 3\%$
$r_3 = 9\%$
$r_4 = 27\%$
$r_5 = 81\%$.

Relatively large changes to the set of rates $\{r_k\}$ make only little difference to the final result.[1] So does changing K to 3 or 5.

A graph of the standard error over the two-year period covered is provided below. The graph would suggest a structural break on 1 June 1988. This is, however, due to the change in data set

1. So the figures have been conveniently rounded.

mentioned above. In fact, up to that date relatively few Bund are used, and a large proportion of the data relates to "runners", i.e. Bund with maturities of 7-10 years, thus spanning the hump in the structure of yields where the fit of the model is distinctly worse.

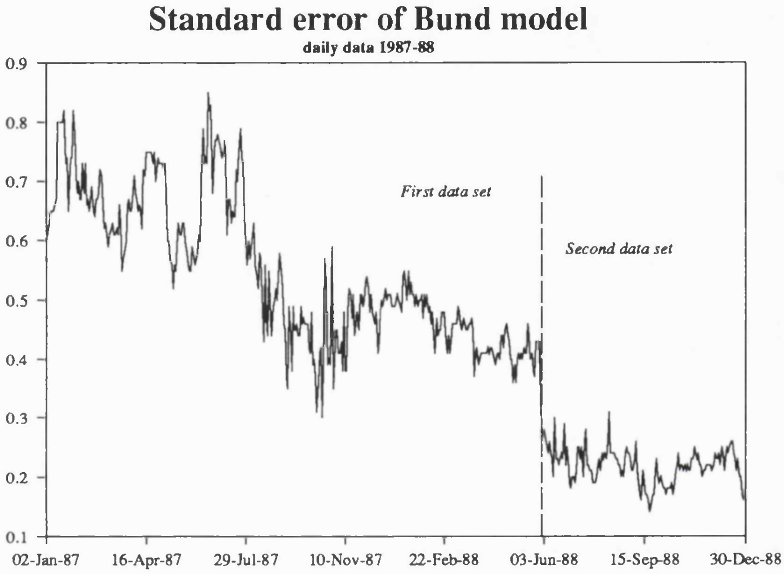

Of greater interest is the gradual improvement in the goodness of the fit throughout the two-year period. This might be partly due to an improvement in the quality of the data. But the more important factor is a reduction in the size of the hump. This has two explanations: firstly, the involvement of foreign investors - which have a higher preference for liquidity - decreased steadily relative to German domestic investors during the two years; secondly, the advent of increasing numbers of traders prepared to undertake yield curve arbitrage and the growing size of the issues involved helped to increase the liquidity in the 8-9 year maturity range.

Estimation results for a day that was fairly average for the period 1 June - 31 December 1988 are given below:

Table 1. Estimation on 1 September 1988

$\beta_1 = -0.23$	$s_1 = 0.043$	$t_1 = -5.2$
$\beta_2 = 0.47$	$s_2 = 0.085$	$t_2 = 5.5$
$\beta_3 = 0.90$	$s_3 = 0.063$	$t_3 = 14.4$
$\beta_4 = -0.13$	$s_4 = 0.027$	$t_4 = -4.8$
$\beta_5 = -0.01$	$s_5 = 0.007$	$t_5 = -1.6$
$\overline{R}^2 = 0.998$	SE = 0.23	
observations = 76	degrees of freedom = 72	

Note: model estimated under the constraint $\beta_5 = \left(1 - \sum_{k=1}^{4} \beta_k\right)$.

The graphs below overlay the par yield curve obtained from the yield surface estimated on 1 September 1988, with the yields of the individual Bund used in the estimation. It is apparent from the graphs and the following table that there is no obvious coupon effect. This is due to the fact that, since issuance occurs almost entirely in the 10-year area, and new issues are always made at coupons such as to ensure prices close to par, the dispersion of coupon levels within a given maturity range is small, and it isn't therefore possible to distinguish a coupon effect from a maturity or liquidity effect. Even in those cases where a coupon effect could emerge, this is swamped by the fact that the stocks with higher coupons usually correspond to the earlier issues which, being less liquid, tend to attract a higher yield.

Bund market yield structure

on 1 September 1988

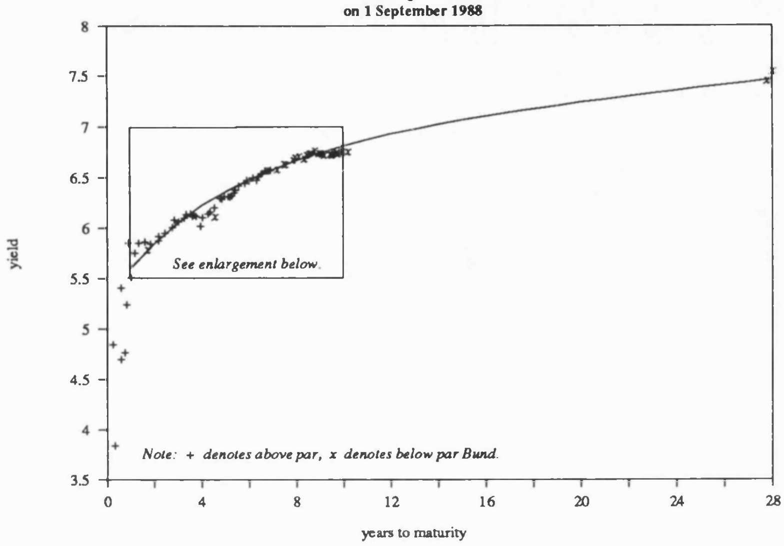

Enlargement

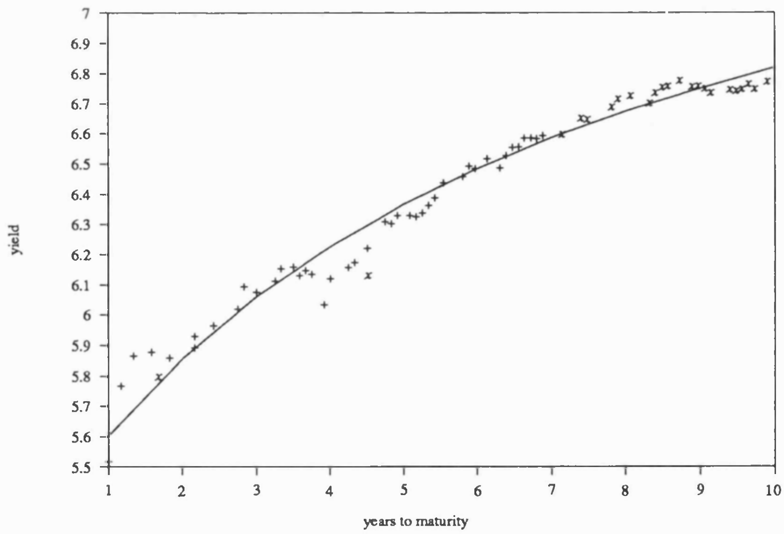

Table 2. Bund list on 1 September 1988

#	stock			price	yield
1	Bund	6.5	1-Dec-88	100.30	4.862
2	Bund	6.75	1-Jan-89	100.80	3.859
3	Bund	10	1-Apr-89	102.30	5.424
4	Bund	7.25	1-Apr-89	101.25	4.716
5	Bund	7.5	1-Jun-89	101.80	4.786
6	Bund	8	1-Jul-89	102.00	5.257
7	Bund	8	1-Aug-89	101.70	5.875
8	Bund	7.5	1-Sep-89	101.75	5.524
9	Bund	7.75	1-Nov-89	102.10	5.772
10	Bund	7.75	1-Jan-90	102.25	5.871
11	Bund	10	1-Apr-90	105.90	5.883
12	Bund	5.75	1-May-90	99.90	5.804
13	Bund	8.25	1-Jul-90	103.95	5.866
14	Bund	7.75	1-Nov-90	103.60	5.897
15	Bund	8.25	1-Nov-90	104.50	5.936
16	Bund	9	1-Feb-91	106.50	5.969
17	Bund	7.5	1-Jun-91	103.55	6.025
18	Bund	10.25	1-Jul-91	110.40	6.100
19	Bund	10.75	1-Sep-91	112.35	6.081
20	Bund	10	1-Dec-91	111.00	6.118
21	Bund	9.75	1-Jan-92	110.40	6.158
22	Bund	9.75	1-Mar-92	110.85	6.164
23	Bund	9.5	1-Apr-92	110.40	6.137
24	Bund	9	1-May-92	109.00	6.152
25	Bund	8.5	1-Jun-92	107.60	6.140
26	Bund	9	1-Aug-92	109.95	6.040
27	Bund	8.75	1-Sep-92	109.00	6.125
28	Bund	7.75	1-Dec-92	105.70	6.162
29	Bund	7.5	1-Jan-93	104.80	6.179
30	Bund	7.5	1-Mar-93	104.80	6.226
31	Bund	6	1-Mar-93	99.45	6.135
32	Bund	8.25	1-Jun-93	107.65	6.314
33	Bund	8	1-Jul-93	106.80	6.308
34	Bund	8.25	1-Aug-93	107.80	6.336
35	Bund	8.25	1-Oct-93	108.05	6.335
36	Bund	8.25	1-Nov-93	108.15	6.332
37	Bund	8.25	1-Dec-93	108.20	6.344
38	Bund	8.25	1-Jan-94	108.15	6.369
39	Bund	8.25	1-Feb-94	108.15	6.394
40	Bund	8	18-Mar-94	107.00	6.443
41	Bund	8.25	20-Jun-94	108.35	6.464
42	Bund	8.25	20-Jul-94	108.30	6.498
43	Bund	8.25	22-Aug-94	108.45	6.490
44	Bund	7.5	20-Oct-94	104.77	6.523
45	Bund	7	20-Dec-94	102.50	6.492
46	Bund	7	20-Jan-95	102.30	6.533
47	Bund	7.25	20-Feb-95	103.45	6.559
48	Bund	7.625	20-Mar-95	105.45	6.563
49	Bund	7.5	20-Apr-95	104.70	6.590
50	Bund	7.25	22-May-95	103.45	6.590
51	Bund	7	20-Jun-95	102.15	6.589
52	Bund	6.75	20-Jul-95	100.80	6.598
53	Bund	6.5	20-Oct-95	99.40	6.603
54	Bund	6.375	22-Jan-96	98.33	6.657
55	Bund	6.375	20-Feb-96	98.33	6.652
56	Bund	5.75	20-Jun-96	94.38	6.693
57	Bund	5.75	22-Jul-96	94.19	6.720
58	Bund	5.5	20-Sep-96	92.53	6.730
59	Bund	6.5	20-Dec-96	98.66	6.706
60	Bund	6.125	20-Jan-97	96.08	6.741
61	Bund	5.75	20-Feb-97	93.58	6.758
62	Bund	6	20-Mar-97	95.13	6.762
63	Bund	5.5	20-May-97	91.73	6.783
64	Bund	6.125	21-Jul-97	95.83	6.761
65	Bund	6.375	20-Aug-97	97.44	6.762
66	Bund	6.75	22-Sep-97	99.94	6.754
67	Bund	6.375	20-Oct-97	97.53	6.741
68	Bund	6.375	20-Jan-98	97.37	6.751
69	Bund	6.25	20-Feb-98	96.52	6.747
70	Bund	6.125	20-Mar-98	95.64	6.753
71	Bund	6	20-Apr-98	94.64	6.768
72	Bund	6.5	20-May-98	98.19	6.753
73	Bund	6.75	20-Jul-98	99.78	6.777
74	Bund	6	20-Oct-98	94.43	6.773
75	Bund	6	20-Jun-16	82.95	7.465
76	Bund	5.625	20-Sep-16	77.70	7.557

2.1.4 Conclusions

The absence of an obvious coupon effect in the German Bund market suggests that a naive yield curve model fitted to Bund yields irrespective of coupon level would for practical purposes provide as good a representation of the market as a discount function model. While this fact is difficult to deny, it is also true that the latter has at least the advantage of generality in that, as illustrated in the following sections, the same formulation of the model appears to be effective in different markets.

In absolute terms, the fit of the Bund model during the period 1 June - 31 December 1988 is good, in the sense that the standard error is usually smaller than the average bid-offer spread.[1] In relative terms, over the same period the Bund model also compares favourably with the gilt model, with a standard error roughly equal to half that obtained on UK gilts.[2] This ought to be expected given the complicated tax effects that exist in the case of gilts. But the fit of the Bund model is not as good as that of the Dutch government bullet model illustrated in the next paragraph. This is probably due to the absence of a "hump effect" in the latter. One may conclude that a weakness of discount function models is probably the difficulty in coping with liquidity effects. This is, however, a problem with virtually all models of bond markets.

1. The bid-offer spread varies from 0.10-0.20 percentage points on the most recent and therefore almost invariably most liquid issues, to 0.20-0.30 percentage points on other liquid issues (the vast majority), to a maximum of 0.50-1.00 percentage points on the two long bonds. This compares with a standard error which is usually of 0.15-0.30 percentage points.

2. The SE on gilts varies in a 0.35-0.55 percentage point range.

2.2 The Dutch government bullet market

2.2.1 Introduction

The Dutch bond market is the oldest in the world, and it was the financial techniques developed in this market that, exported to England at the time of the Glorious Revolution, formed the basis of the modern day UK gilt market. However, by contemporary standards, the Dutch bond market is on the whole antiquated, generally lacking in liquidity and depth. As a result, only a small proportion of it is truly liquid and internationally traded.

The most liquid Dutch bonds are those issued by the government, which account for roughly two thirds of the total volume outstanding. The vast majority of Dutch bonds in existence include sinking fund arrangements[1] with or without additional early redemption or extension options.[2] Indeed, until 1 January 1986 this was the only type of bond allowed for official listing.

Sinking fund bonds, while very popular with domestic issuers and investors alike, have a number of disadvantages. The relative

---

1. Bonds issued with sinking fund arrangements are redeemed by instalments on a number of dates (usually spanning several years) according to a pre-established schedule. The bonds to be redeemed on any given date are, with the exception of the final instalment, selected by random drawing which takes place shortly before the partial redemption date. Therefore, the owner of a sufficiently large holding can expect his bonds to be redeemed in proportion to the amounts scheduled for each drawing; the same applies to the owner of bonds held in a large pool, as is the case when the custodian is one of a number of settlement agencies. Since a large proportion of investors fall into one of these two categories, it might be sensible to apply a discount function model to the expected cash flows, effectively treating a sinking fund bond as a basket of bullets with a nominal value equal to a fraction of par.

2. It is common in the Dutch bond market for sinking fund bonds to include provisions whereby the government has the option to redeem the bonds outstanding on certain dates at pre-established prices (callable bonds), or the holder has the option to convert on certain dates the bonds into other bonds of longer maturity (extendible bonds).

complexity of their cash flow pattern ensures that it is difficult to make comparisons of value between different issues, and also makes it more awkward to use them in conjunction with interest rate or currency swaps.[1] Furthermore, when redemptions start, the size of the issue shrinks rapidly, and therefore sinking fund bonds tend to become very illiquid. For these and other reasons they are generally disliked by international investors, who require a higher yield to be induced to buy them.

With effect from 1 January 1986 the Dutch government decided to lift the existing restrictions on bullet bonds, and immediately started to issue bullets on a large scale. These issues have proved a great success, especially with foreign investors. The Dutch government bullets, of which over 20 issues are now in existence, are by far the most liquid Dutch bonds. Whereas the overall size of the Dutch government bullet market is very small, the individual issues are rather large, usually around Dfl 3-5bn, and this enhances their liquidity. However, the spectrum of maturities available is narrow, ranging only from 4 to 10 years. A significant proportion of trading in these bonds is done in London, which makes the market directly comparable with the other markets examined in this study.

Despite their dislike for sinking fund bonds, foreign investors have always maintained a significant presence in the Dutch bond market. This is largely due to the favourable tax treatment[2] and, especially in recent years, to the hard currency status of the Guilder. When the Dutch government bullet market was born it

1. In brief, these are agreements whereby cash flows are "swapped" for other cash flows determined e.g. by a variable interest rate as opposed to a fixed one (interest rate swap), or denominated in a different currency (currency swap). These transactions represent an extremely large and growing volume of business in the more advanced financial markets.

2. Dutch bonds have never been subject to withholding tax and occasional rumours that a tax might be introduced have never been taken very seriously.

was from the start dominated by foreign investors, and as such it is one of the best examples of a bond market with virtually no tax effects. In this it contrasts strongly with the sinking fund segment of the bond market, which is affected in much the same way as the UK gilt market[1] by the presence of domestic investors with a strong tax-induced preference for capital gain relative to income.

1. See chapter III.

2.2.2 Empirical results

The Dutch government bullet model was estimated using daily data covering the whole of 1988 obtained from Bloomberg. Because the market was so new there were only 11 bonds in existence at the beginning of the year, whereas by the end of the year there were 19. All available bond prices were used in the estimation.

As illustrated in the graph below, the model standard error appears to decline substantially throughout the sample period. This is probably due to the fact that, as the size of the market increased from its initial minuteness, the liquidity and pricing accuracy improved.

Standard error of Dutch gov't bullet model

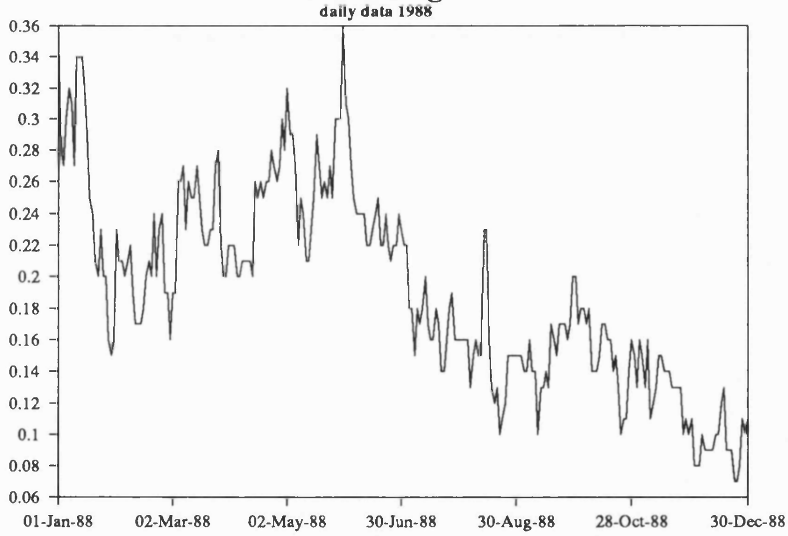

Throughout, the standard error is almost always within the bid-offer spread.[1] As indicated previously, the fit also compares favourably with that of the Bund model.[2]

The discount function formulation used for Bund was found to be adequate for Dutch bullets and so was kept unchanged. Again, it was found that relatively large changes to the set of rates $\{r_k\}$, including changing K, make little difference to the fit. Estimation results for 1 September 1988, which are reasonably representative of the performance of the model, are given below:

Table 3. Estimation on 1 September 1988

$\beta_1 = 1.72$	$s_1 = 0.82$	$t_1 = 2.1$
$\beta_2 = -2.76$	$s_2 = 1.38$	$t_2 = -2.0$
$\beta_3 = 2.53$	$s_3 = 0.71$	$t_3 = 3.6$
$\beta_4 = -0.56$	$s_4 = 0.18$	$t_4 = -3.2$
$\beta_5 = 0.06$	$s_5 = 0.03$	$t_5 = 2.4$
$\overline{R}^2 = 0.984$	SE = 0.15	
observations = 17	degrees of freedom = 13	

Note: model estimated under the constraint $\beta_5 = \left(1 - \sum_{k=1}^{4} \beta_k\right)$.

1. This is usually 0.20-0.30 percentage points, compared to a standard error ranging between 0.10 and 0.30 percentage points most of the time.

2. The average standard error of 0.19 is about 15% lower than that obtained on Bund over the second sample period.

The graphs below overlay the par yield curve calculated from the estimated yield surface on 1 September 1988 with the yields of the individual bonds in existence at the time and used in the estimation. From the graphs and the following table it is apparent that a coupon effect may exist, although it is not clear cut. This is probably due to the fact that the dispersion of coupons is very small and, in view of the relative flatness of the yield structure, the predicted effect would in any case be extremely small.[1]

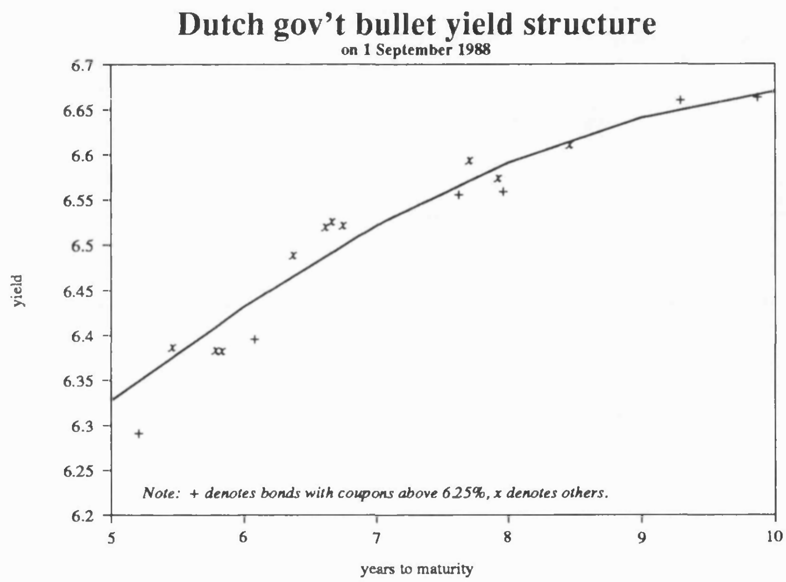

1. A 1% point difference in coupon at a maturity of 10 years should, according to the model, be worth less than 0.01% in yield terms.

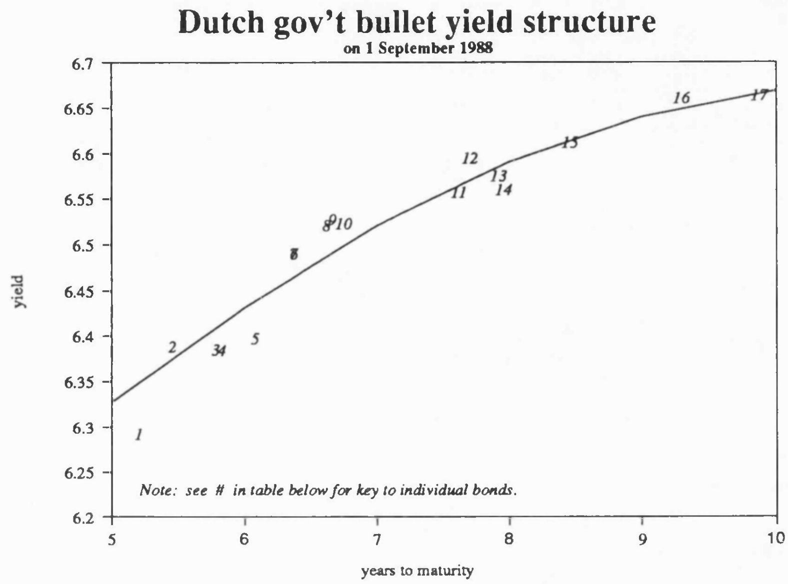

Table 4. Dutch gov't bullet list on 1 September 1988					
#	**stock**			**price**	**yield**
1	Nether	7	15-Nov-93	103.00	6.293
2	Nether	6.25	15-Feb-94	99.30	6.389
3	Nether	6	15-Jun-94	98.15	6.386
4	Nether	6	1-Jul-94	98.15	6.385
5	Nether	6.5	1-Oct-94	100.45	6.398
6	Nether	6.25	15-Jan-95 'A'	98.70	6.491
7	Nether	6.25	15-Jan-95 'B'	98.70	6.491
8	Nether	6	15-Apr-95	97.20	6.522
9	Nether	6.25	1-May-95	98.50	6.528
10	Nether	6.25	1-Jun-95	98.50	6.524
11	Nether	6.5	15-Apr-96	99.60	6.558
12	Nether	6	15-May-96	96.45	6.596
13	Nether	6.25	1-Aug-96	98.00	6.576
14	Nether	6.5	15-Aug-96	99.60	6.561
15	Nether	6.25	15-Feb-97	97.60	6.613
16	Nether	6.375	15-Dec-97	98.00	6.662
17	Nether	6.5	15-Jul-98	98.80	6.665

2.2.3 Conclusions

The fact that the same model used for Bund can be applied to Dutch bullets to yield such sensible results confirms the usefulness of this type of approach. This remains true despite the observation that one of the main theoretical advantages of using discount functions compared to yield based methodologies, i.e. the accurate measurement of the duration effect, is in practice often not very important for the purpose of describing the yield structure. The accurate measurement of the duration effect is, however, generally much more important for pricing purposes. This is illustrated in the graphs below.

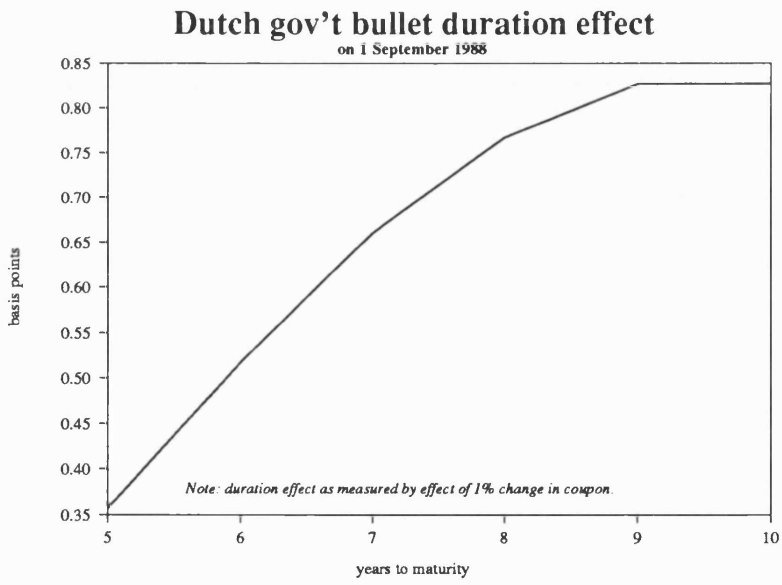

Price equivalent of duration effect

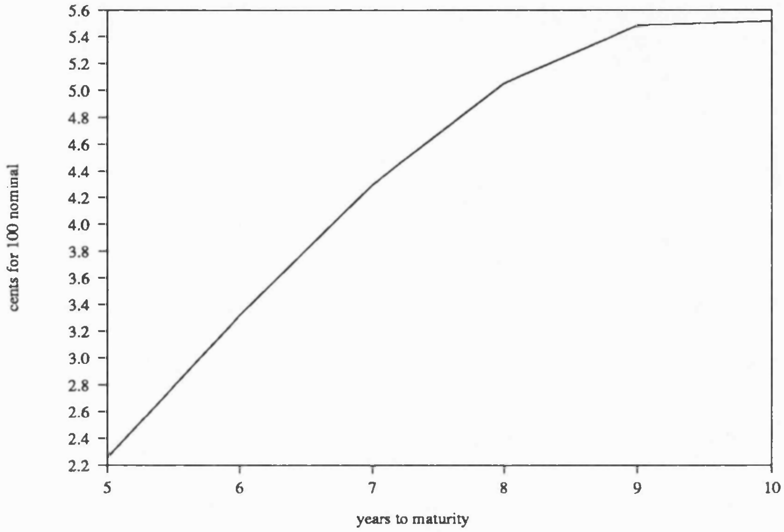

It would appear that discount functions really come into their own for picking up coupon effects only when the dispersion of coupons is great and, possibly, when taxation is the main cause. This is illustrated in the following study of the UK gilt market.

III. THE UK GILT MARKET

3.1 Introduction

UK gilts or gilt-edged stocks are sterling-denominated bonds issued on the domestic market by the UK Treasury or bearing the guarantee of the UK government. The gilt-edged market is the third largest government bond market in the world[1] and is one of the most liquid. There are two main categories of gilts: "conventional gilts" and "index-linked gilts". The distinguishing factor is that conventional gilts have fixed coupons and are redeemed at par, whereas in the case of index-linked gilts the coupons and the principal repayment are indexed to the RPI. Conventional gilts form the bulk of the market, with about 100 different stocks in existence.[2] There are relatively few index-linked gilts, since they were first issued only in 1981 and they have not yet been relied upon much by the Treasury as a source of finance.[3] Because of their special features, index-linked gilts trade on a separate market from conventional gilts, and often behave quite independently of the latter. The remainder of this analysis will focus on conventional gilts.

Most conventional gilts are redeemed within a fixed time period, i.e. are "dated". However, six of them are "undated", i.e. can be redeemed at the Treasury's request after certain fixed dates, with no other time

1. After the US and Japanese markets. But if bills were included it would be overtaken by the Italian market.

2. At the time of writing, this number is rapidly declining because of the absence of new issues.

3. This is partly due to the fact that the market for index-linked gilts took time to develop, and partly that the sharp decline and then reversal in the public sector borrowing requirement made it unnecessary to issue large amounts of stock.

limit.[1] In practice, however, since the coupons on these gilts are very low, it is unlikely that the Treasury will wish to redeem them in the foreseeable future, and therefore undated gilts are treated by the market as though they were irredeemable. Because they were issued long ago and in current nominal terms have a relatively small size, they are very illiquid, and tend to appeal to a small subset of investors. In the rest of this study undated gilts will be ignored.

Conventional dated gilts are either redeemed on a fixed date, or at the Treasury's request within two specified dates. These "double-dated" gilts include therefore an option in favour of the Treasury. If the option is "at the money", i.e. the price of the gilt is close to par, the double-dated gilt should be expected to yield more than a comparable single-dated gilt, owing to the negative value of the option. In the gilt model, double-dated gilts that contain a substantial option value element are neglected, whereas double-dated gilts that are far "out of the money" are treated as single-dated gilts of the appropriate maturity.

Finally, certain gilts include an option in favour of the holder to convert on specified dates and at pre-established rates the gilt held into one or more others that may be created as a result of the conversion option being exercised. These "convertible gilts", of which few are in existence at any one time, are used in the gilt model only when the last conversion date has expired.

1. But one of these gilts incorporates a sinking fund arrangement.

3.2 The participants

3.2.1 Investors

The table below gives a break-down of the gilt market according to type of holder and maturity band:

Table 5. Nominal gilt holdings on 31 March 1988

percentage break-down by holder and maturity

years' life:	<5	5-15	>15
pension funds	0.9	9.9	10.6
overseas holders	5.2	6.2	0.9
personal sector	4.5	4.2	2.1
life assurance	1.2	12.6	11.1
general insurance	1.4	2.7	0.3
monetary sector	1.5	2.6	2.0
building societies	5.4	0.6	0.0
others	7.9	5.6	0.3
total	28.0	44.6	27.4

Sources: CSO, Bank of England, my estimates.

It emerges that the largest investors in gilts are life assurance companies and pension funds. These dominate the medium and long maturities, with holdings amounting to almost two thirds of the total outstanding in those categories. Both groups have a preference for longer maturities since these reduce their overall portfolio risk in view of the extended time-profile of their liabilities. For tax reasons that are explained below life assurance companies have a relative preference for low-coupon gilts while pension funds prefer higher coupons.

The rest of the market, neglecting the residual, is almost equally divided amongst overseas sector, personal sector, and banks, building societies and insurance companies. The overseas sector holds mainly high-coupon short and medium maturities, especially FOTRA gilts.[1] The personal sector dominates the low-coupon gilts at short and medium maturities. Banks, building societies and insurance companies invest mainly in the medium and high-coupon short maturities.

These observations based on statistical data and anecdotal evidence are substantiated by the gilt model estimation results, which are based on endogenous segmentation and are therefore independently derived.

1. FOTRA is an acronym for "free of tax to residents abroad". A description of FOTRA gilts is given in section 3.4.1.

3.2.2 Dealers

On 27 October 1986, date of the so-called Big Bang,[1] the current structure of the gilt market came into being. The new market is based on gilt edged "market makers", i.e. dealers who have undertaken to quote continuous and realistic bid and offer prices to non market makers in a suitable number of gilts in appropriate size, and have been recognized as such by the Bank of England who constantly monitors and supervises their activity.

In exchange for this obligation, market makers have a number of privileges. They can display firm bids and offers to other market makers through the inter-dealer brokers' screen-based systems, which protect their anonymity. They have direct access to the Bank of England's dealing room. They can borrow gilts through the money brokers and can thus "sell short", i.e. sell gilts that they do not own.[2]

In addition to the market makers there are a few brokers who act as intermediaries between investors and market makers in an advisory capacity to the former.

---

1. The long-awaited reform of the Stock Exchange rules to comply with the Restrictive Practices Act.

2. This capability is not cheap, since it costs usually 0.5-1% per annum. Also, there are many gilts that it is virtually impossible to borrow at any price. Typically, these are the low-coupon gilts that it would usually, for tax reasons, be most advantageous for market makers to sell short. This point is explained in section 3.5.3.

3.3 The market place

The increased competition that followed the abolition of restrictive practices with Big Bang caused a substantial reduction in dealing costs for investors. Bid and offer prices quoted by market makers now are net, i.e. they are inclusive of all charges to the purchaser. The bid-offer spread therefore represents the greater part of the total cost of dealing[1] and provides a good measure of the latter. Bid-offer spreads on gilts vary nowadays from less than 1/32 on gilts of very short maturities, to 2-4/32 on liquid gilts of longer maturities, to 6-8/32 on less liquid gilts especially at longer maturities, and to possibly 16/32 on very illiquid issues. Wider spreads would apply in the case of abnormal sizes.

Since 10 February 1986 all gilt prices are quoted on a "clean" basis. A clean price is equal to the total price payable less the "accrued interest". The accrued interest is given by the formula:[2]

$$a = c\,[\delta/365]$$

where c is the coupon and δ is the "days of accrual", that is the number of days between the last coupon payment date and settlement date[3] if the gilt is traded "cum-dividend", i.e. inclusive of the next coupon payment, or minus the number of days between settlement date and the next coupon payment date if the gilt is traded "ex-dividend". Gilts are traded cum-dividend or ex-dividend according to the ex-dividend calendar published by the Bank of England. All the prices used in the UK gilt model estimation are for ordinary bargains, and the standard ex-dividend and settlement rules are therefore applied.

1. The investor will, of course, have to bear his own internal transaction costs and maybe external costs relating to settlement. Purchases of gilts are exempt from stamp duty.

2. A more complicated formula applies in the case of partly paid gilts.

3. Ordinary settlement is on the next business day, except in the case of an auction stock which is still trading on a "when issued" basis.

3.4 Taxation of gilts

3.4.1 Income & capital gain

Holders of gilts are subject to taxation of income and of capital gain. The amount of taxable income is calculated on an accrual basis, whereas the relevant date is usually given by the payment date of the coupon to which the accrual relates. Before 28 February 1986 taxable income was not calculated on an accrual basis, and therefore income could be converted into capital gain by selling cum dividend and buying ex; this used to give rise to "bond washing", whereby investors could convert income into capital gain and vice versa to take advantage of the different rates of tax. The only modest form of bond washing now available relates to tax-timing considerations, whereby investors who are subject to withholding tax have an incentive not to receive coupons[1] and, provided that they are planning not to receive the coupons, to hold gilts with coupon payment dates at the beginning of their tax year.

As a rule, coupons are paid by the Bank of England net of the basic rate of income tax. In certain circumstances, usually when the investor is exempt from tax on income or subject to a reduced rate, the coupon may be paid gross. These special cases can be summarized as follows:

- certain gilts (denominated FOTRA) may be exempted upon request by the holder if he qualifies as a non-resident[2]

1. This is known as coupon washing.

2. For tax purposes non-residents are defined as individuals not ordinarily resident in the UK and companies managed and controlled outside the UK.

- gilts held by approved pension funds may be exempted upon request

- gilts held by entities enjoying sovereign immunity may be exempted upon request

- gilt holdings falling within the purview of an agreement to this effect between the UK and a foreign government may be exempted or subject to a lower rate of tax in accordance with the agreement.

The amount of taxable capital gain is calculated, consistently with the taxation of income, as the difference between sale (or redemption) and purchase price, net of accrued income. The relevant date for tax purposes is usually given by the date of disposal (or redemption).

3.4.2 Taxation

Income and capital gains are subject to income tax or corporation tax. As from 2 July 1986 capital gains are no longer subject to capital gains tax, while remaining subject to corporation tax when they represent a component of a company's ordinary trading profit. Investors in gilts may be broadly grouped into three categories, according to their tax status: gross, net and net net.

A. Gross investors

These don't pay tax on either their income or their capital gains. Gross investors may be classified as follows:

A1. Approved pension funds and related business

A2. Residents abroad

These constitute rather a heterogeneous grouping, since double-taxation agreements may apply to them, and they may be subject to tax at various rates in their country of residence.[1] With regard to UK taxation, while exempt from tax on FOTRA gilts, they are subject to withholding tax on non-FOTRA gilts unless they are covered by a double-taxation agreement - and then only within the terms of the agreement - and in this case they therefore usually wash the coupon, i.e. they sell before the next ex-dividend date, thus avoiding the tax.

---

1. In a sense, they are therefore unclassifiable. However, in practice the majority of foreign investors behave as though they did not pay tax.

B. Net investors

These pay tax on income but not on capital gains. They may be grouped as follows:

B1. Life funds and most non-financial companies

These are subject to corporation tax on income only, not on capital gains.

B2. Personal sector and most investment companies

These are subject to income tax, or corporation tax and income tax on income only.

C. Net net investors

These are taxed on income and capital gains. The major representatives of this group are:

C1. Banks, building societies, insurance companies and most other financial companies

These pay corporation tax on income and capital gains.

The portion of income and capital gains that represents the company's ordinary trading profit is subject to corporation tax even in the case of gross or net funds. This doesn't significantly alter the broad classification given above. However, during the sample period used in the gilt model estimation, life assurance and insurance companies that carried out a mixture of business subject to different regimes were taxed on the aggregate at the average rate applicable, rather than separately on the individual types of business. Therefore, most of these companies did not fall into any of the tax brackets above.

3.4.3 Rates of tax[1]

The rates of income tax applicable to the sample period are as follows:

Table 6. Rates of income tax

year	minimum	maximum
86/87	29%	60%
87/88	27%	60%
88/89	25%	40%

The rates of corporation tax applicable are as follows:

Table 7. Rates of corporation tax

year	small companies[2]	standard
86/87	29%	35%
87/88	27%	35%
88/89	25%	35%

1. It may occur that an investor of a type that would usually be taxed enjoys total or partial relief because of special circumstances, e.g. because his personal income is sufficiently low. It is assumed that these cases carry negligible weight in gilt pricing.

2. Small companies that are not small enough to qualify for the lower rate may nonetheless enjoy partial tax relief.

3.4.4 Conclusion

From the point of view of UK taxation one may distinguish three main types of investor. Gross investors, who pay no tax, and dominate the medium to high-coupon gilts of medium to long maturities. Net investors, who pay tax on income - but not on capital gains - at rates of up to 60% in 86/87 and 87/88, and of up to 40% in 88/89 and dominate the low to medium coupon gilts of medium to long maturities in the case of life funds[1] and of short to medium maturities in the case of the personal sector. Net net investors, who pay tax on income and capital gains at rates of 29% to 35% in 86/87, of 27% to 35% in 87/88, and of 25% to 35% in 88/89, and dominate the medium to high-coupon gilts of short maturities.

1. Note, however, the point made about life assurance companies in section 3.4.2.

3.5 Gilt modelling literature

The UK gilt market has probably been modelled more than any other bond market over the past twenty years. This is largely due to the fact that the size of the market, the number of issues and the tax regulations[1] have, for many years, maintained a liquid market in a large number of issues spanning a wide range of coupons. Thus, models have been developed both for practical purposes, such as for fairly pricing new issues or detecting arbitrage opportunities, and for theoretical purposes, such as exploring the workings of a segmented market.

A "golden age" of gilt market analysis is probably now over, since the government's buying-in policy and the lack of new issuance have dramatically reduced the size of many stocks and the number outstanding, thus concentrating liquidity in relatively few stocks and reducing the importance of the market as a whole. The following subsections provide a brief guide to the best known gilt models developed to date.

1. Which have generally encouraged stock switching activity.

3.5.1 Bank of England model

The original Bank of England gilt model was based on the work of J. P. Burman and W. R. White, details of which are given in their article in the Bank of England Quarterly Bulletin [6], and underwent numerous minor modifications during the following fifteen years, as described in a series of articles [1, 3, 4, 5], before being supplanted by a new approach - which is of little if any theoretical interest - as explained in a further article in the Bank of England Quarterly Bulletin [2]. With the benefit of hindsight the most obvious difficulty with the original model was the constant need to adjust the specification. This is largely due to some ad-hoc features of the model, that tend to restrict its generality.

In brief, the Burman and White approach consists in assuming that the gilt market is divided into two segments, one for shorter maturities and one for longer maturities; in each of these a representative investor expects to receive on all stocks a net holding yield up to a planning horizon, as determined by a prospective net redemption yield at the horizon; tax rates, expected net holding yield, planning horizon and prospective net redemption yield are assumed to be equal across stocks but different in each segment. The two price functions obtained are then spliced together to produce the market price function. Subsequent changes to the model affected the shape of the price-coupon relationship, which was changed from linear to convex,[1] and the splicing mechanism.

1. This is explained in the next section.

3.5.2 Clarkson model

This approach is laid out in detail in Clarkson [13]. The main achievement of the Clarkson model was to emphasize the importance of segmentation. In models without segmentation the relationship between price and coupon is generally linear.[1] For example, in the case of a flat net yield curve and a unique constant rate of coupon tax, the price of a bond on a coupon payment date would be given by:

$$p = c \sum_{i=1}^{n} (1 - \tau/100)(1 + y/100)^{t_i} + 100(1 + y/100)^{t_n}$$

where p denotes the price, c the coupon, y the (annual) yield, τ the rate of tax, t_i the time in years to the i-th cash flow and n the number of cash flows.

Clarkson recognized that if investors paying different rates of tax were involved, the relationship between price and coupon could easily become convex.[2] In fact, if selling short is not allowed, given a set I of investors, each of which is endowed in equilibrium with a linear indifference curve $p_i(c)$ as described by a price equation of the type stated above, the actual bond price must lie on the lower boundary of the set $\{(p,c) \in \Re^2 \mid p \geq p_i(c), \forall i \in I\}$ which is convex, being the intersection of (weakly) convex sets.

1. This type of result is very general. The US is an exception, since a kink occurs at par due to the option to amortise capital losses. The gilt model discussed in this study is linear within each market segment.

2. As explained in chapter I, this can occur even without differential taxation, but is then less likely.

3.5.3 Schaefer model

Schaefer's approach is similar to Clarkson's in that it allows for convexity of the price-coupon relationship, but differs in that it incorporates this feature into a formal discount function framework. However, rather than attempting to model the market by taking a small number of investor types and assuming segmentation according to a representative investor for each type, as is done in this study, his main interest is to illustrate how segmentation occurs as a result of rational portfolio behaviour on the part of investors faced by different tax schedules. The basic assumptions of the Schaefer approach are, however, very similar to those adopted in this study, and are therefore discussed at length in other sections.

The reason why Schaefer introduces segmentation deserves separate mention. In a complete bond market with differential taxation, unless short and long positions are treated asymmetrically there is no equilibrium because of the possibility of unlimited tax arbitrage.[1] For example, if π denotes the relative price of a cash flow subject to income tax in terms of a cash flow that is tax free, τ_i denotes the rate of income tax applicable to the i-th individual and I denotes the set of individuals, equilibrium requires:

$$\forall i \in I : \pi = (1 - \tau_i/100).$$

Otherwise, two individuals with, say, rates τ_1 and τ_2 with $\tau_1 > \tau_2$, could trade with mutual advantage and to the detriment of the tax collector at any price $\pi \in (\pi_1, \pi_2)$, where $\pi_i = (1 - \tau_i/100)$.

1. At a practical level, a gilt-edged market maker, who can borrow gilts and thus sell short, would arbitrage between low and high coupon gilts were he not constrained by the cost of borrowing.

3.6 A discount function model with taxation

3.6.1 One investor type case

Let α and γ be the rates of tax on income and capital gain applicable to a representative investor. Let $n(p, b, \alpha, \gamma)$ be the vector of after-tax cash flows on a bond - assuming it is bought and held to maturity - given the pre-tax flows b, the price p, and the tax rates. In a discount function framework the price function for such a bond is implicitly defined by:

$$p = n'v.$$

In standard notation, the model equation is then of the form:

$$p = p(v, b, \alpha, \gamma) + \varepsilon.$$

Using net of tax cash flows presents a major difficulty, because the timing of the tax payments depends in practice on a number of circumstances, some of which are unknown. The most serious problem is probably the timing of tax payments relating to capital gains.[1] Since most tax codes require the tax to be paid after the gains are realized, the bond holder is usually given considerable control over when he pays the tax. The holder of a bond may, for example, wish to realize his losses immediately in order to benefit from the tax reduction, but to defer the realization of capital gains. There exists therefore in this respect a "tax-timing option". The buy-and-hold assumption made here is clearly inappropriate, but difficult to avoid. This point is discussed further in chapter V.

---

1. The situation in the gilt market is relatively straightforward compared to the US bond markets, where a distinction has to be made between short and long-term capital gains.

3.6.2 Several investor type case

The approach followed in this and the next chapter is to assume segmentation, whereby each bond is bought, and therefore priced, according to the type of investor that puts the highest value upon it. As already mentioned, this model requires that long and short positions are treated asymmetrically. The asymmetry is obtained here, as in the case of section 1.3.2, by assuming that short positions are not allowed.

The model would then be of the form:

$$p = \max(\{p(v_i, b, \alpha_i, \gamma_i), i \in I\}) + \varepsilon$$

where I is the set of investor types. In practice, in order to obtain a smooth and well-behaved function, a generalisation of this model is used, based on the following result.

Lemma

Consider the finite set Z containing positive real numbers. Let $\pi(Z,r)$ be the power mean of Z of order r, i.e. in standard notation:

$$\pi = \begin{cases} \left(\sum_{z_i \in Z} z_i^r \right)^{1/r} & \text{if} \quad r \neq 0 \\ \left(\prod_{z_i \in Z} z_i \right)^{1/\aleph(Z)} & \text{if} \quad r = 0 \end{cases}$$

where $\aleph(Z)$ denotes the cardinality of Z.

Then:

(i) $\min(Z) \leq \pi \leq \max(Z)$

(ii) $r_1 > r_2 \Rightarrow \pi(Z,r_1) \geq \pi(Z,r_2)$

(iii) $\lim_{r \to -\infty} \pi(Z,r) = \min(Z)$

(iv) $\lim_{r \to +\infty} \pi(Z,r) = \max(Z)$.

The equation used for modelling the gilt market is of the form:

$$p = \pi(\{p(v_i, b, \alpha_i, \gamma_i), i \in I\}, r) + \varepsilon$$

where r is assumed to be constant over time, and is estimated to be very large (300-500) but finite.

3.7 The gilt model

Let p be the (clean) price of a gilt, a be the accrued, t_a be the time to the next coupon date, c be the vector of before tax coupon payments, α and γ be the rates of tax, v be the discount function and t the time to redemption. Then, in standard notation, and neglecting for simplicity the case of partly paid stocks,[1] the price function for a representative investor is implicitly defined by:

$$p = (1-\alpha)\, c'v + \alpha\, a\, v_{t_a} + p\, v_t + (1-\gamma)(100-p)\, v_t - a$$

or equivalently by:

$$p = (1-\alpha)\, c'v - (1-\alpha v_{t_a})\, a + 100\,(1-\gamma)\, v_t + \gamma\, p\, v_t.$$

So the price function for the representative investor is given by:

$$p = \{(1-\alpha)\, c'v - (1-\alpha v_{t_a})\, a + 100\,(1-\gamma)\, v_t\}/\{1-\gamma v_t\}.$$

In accordance with the description given of the tax regime of gilts it is assumed that the model may be built around three representative investors, one for each type of tax status: gross, net and net net.

The relative tax rates are respectively:

(i) $\alpha_1 = 0, \gamma_1 = 0$

(ii) $\alpha_2 = \text{unknown}, \gamma_2 = 0$

(iii) $\alpha_3 = \text{unknown}, \gamma_3 = \alpha_3.$

1. In the case of a new issue, part of the issue price payment may be deferred, so that one or more negative cash flows must be included in the price function.

By a change in notation, letting $\alpha_2 = \alpha$ and $\alpha_3 = \gamma$, one gets the final model equation:

$$p = \pi(\{p_1(b, v_1), p_2(b, \alpha, v_2), p_3(b, \gamma, v_3)\}, r) + \varepsilon$$

where $p_1(.,.)$, $p_2(.,.,.)$, $p_3(.,.,.)$ denote the price functions for gross, net and net net representative investors respectively.

3.8 Gilt yield surface

Allowing for taxation fundamentally affects the shape of a yield surface. This is illustrated by the following three graphs, describing yield surfaces for individual investor types generated by the same par yield curve under the ad hoc assumption that the curve lies on each surface.[1] The three surfaces correspond, respectively, to the assumptions:

(i) $\alpha_1 = 0, \gamma_1 = 0$

(ii) $\alpha_2 = 35\%, \gamma_2 = 0$

(iii) $\alpha_3 = 35\%, \gamma_3 = 35\%.$

Yield surface (i)

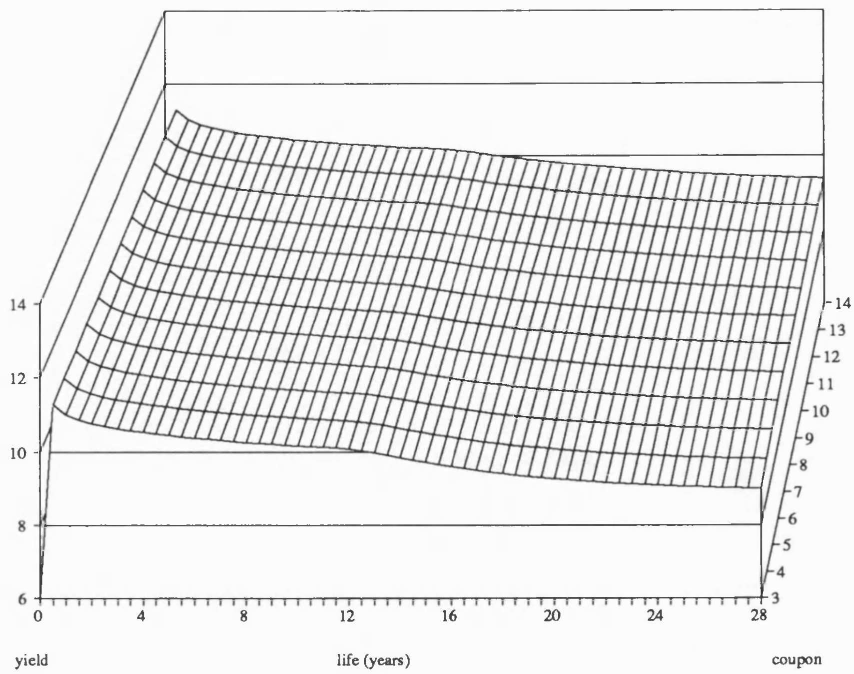

1. The actual estimated par yield curve for 1 September 1988 is used in these examples.

Yield surface (ii)

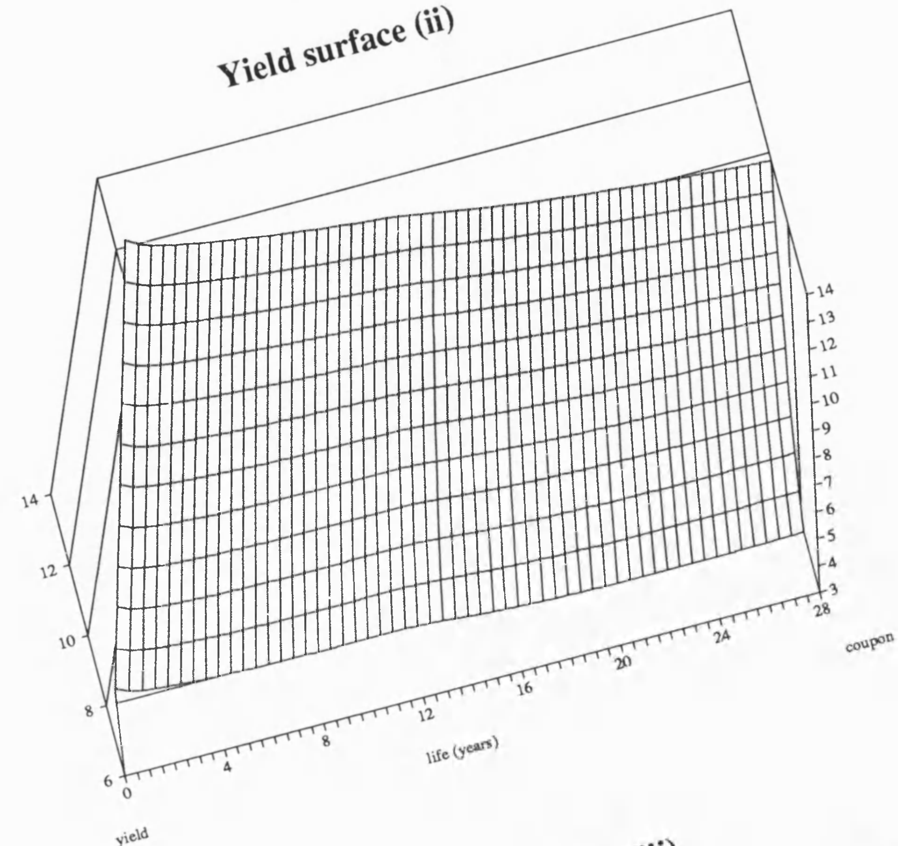

Yield surface (iii)

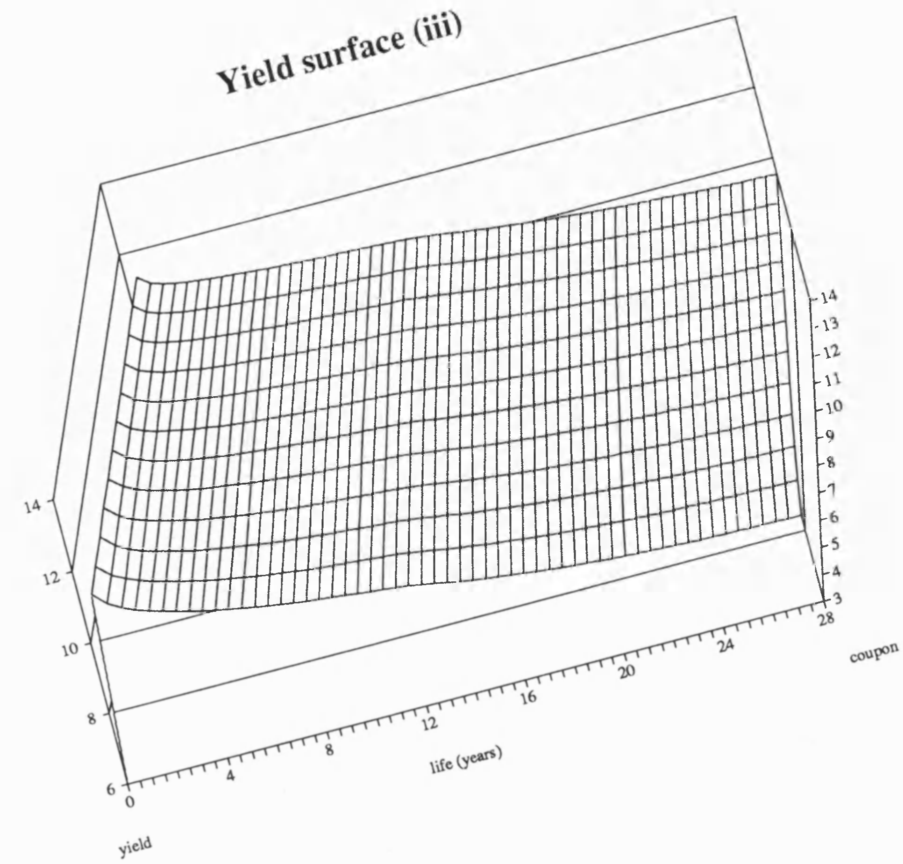

The first graph provides an example of a pure duration effect. This is absent at the shortest maturity, and rises at longer maturities to a level which, despite the fact that the par yield curve is quite steep, is relatively modest.[1]

The subsequent two graphs show how tax effects can easily swamp duration effects. The graphs also illustrate how the nature of the effects is quite different under the two taxed regimes. Under the net regime the effect reaches its maximum at the shortest maturity, and decreases throughout maturities. Under a net net regime the effect is, as in the case of a pure duration effect, absent at the shortest maturity, and rises to a relatively modest level at the longer maturities. In this latter case the behaviour is due to the advantage capital gain enjoys compared to income, owing to the delay in payment of the tax. The existence of the tax-timing option should, therefore, be expected to reduce the effect at higher coupon levels, thus causing the slope of the surface in the direction of the coupon to flatten and possibly even, depending on the type of model assumed, to become negative.

Given the very different types of effect under different regimes, one might expect it to be easy to distinguish between them on empirical grounds. However, the simultaneous presence of a number of different regimes can complicate matters considerably.

1. A change in coupon of 1% at long maturities corresponds to only about 4 basis points in yield. This is, however, much more than that observed in the case of the Dutch bullet market (see section 2.2.3).

The yield surface corresponding to the UK gilt model is a composite surface approximately obtained by joining up the minima of the actual surfaces applicable to each investor type.[1] Estimated individual investor type surfaces corresponding to the same three tax assumptions as before are illustrated in the following three graphs. Comparing these and the previous ones gives an idea of the magnitude of the bias that could derive from neglecting market segmentation.

Yield surface (i')

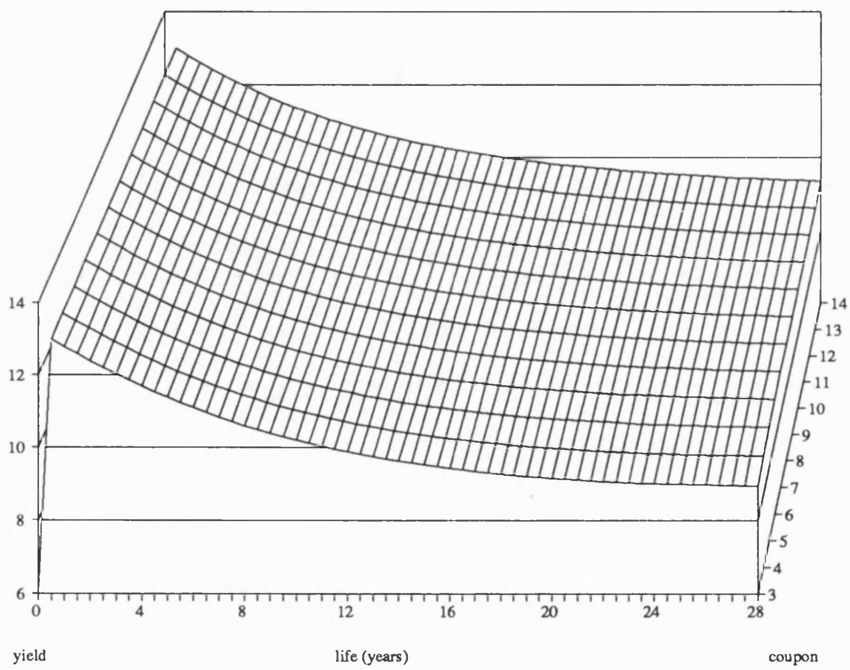

1. These surfaces are different from the previous ones in that they are the actual estimated surfaces rather than the theoretical ones generated - under different ad hoc single investor type assumptions - by a given par yield curve.

Yield surface (ii')

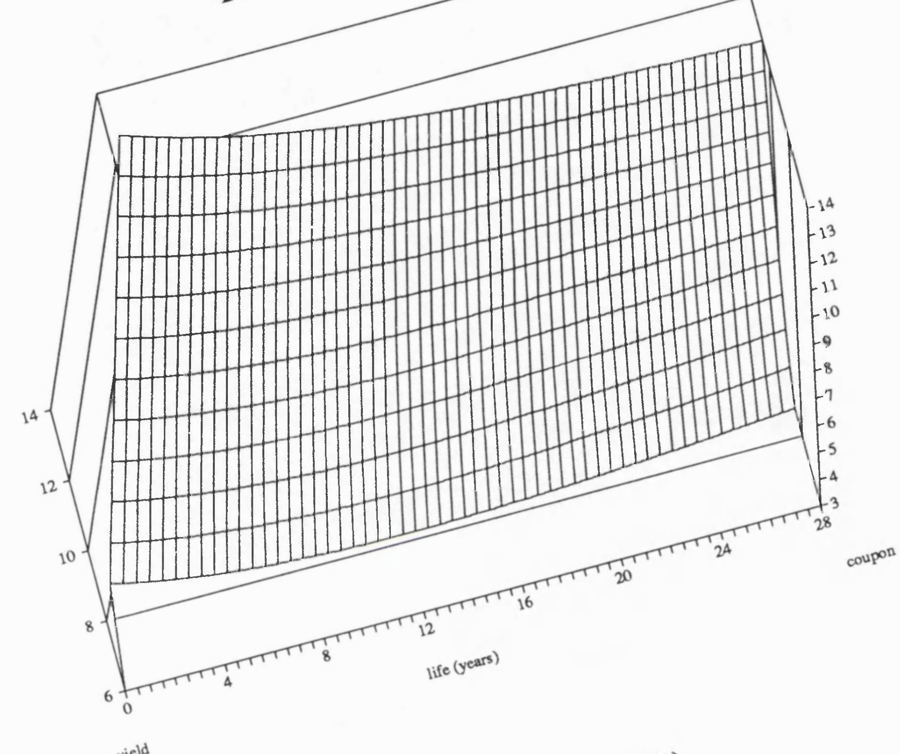

Yield surface (iii')

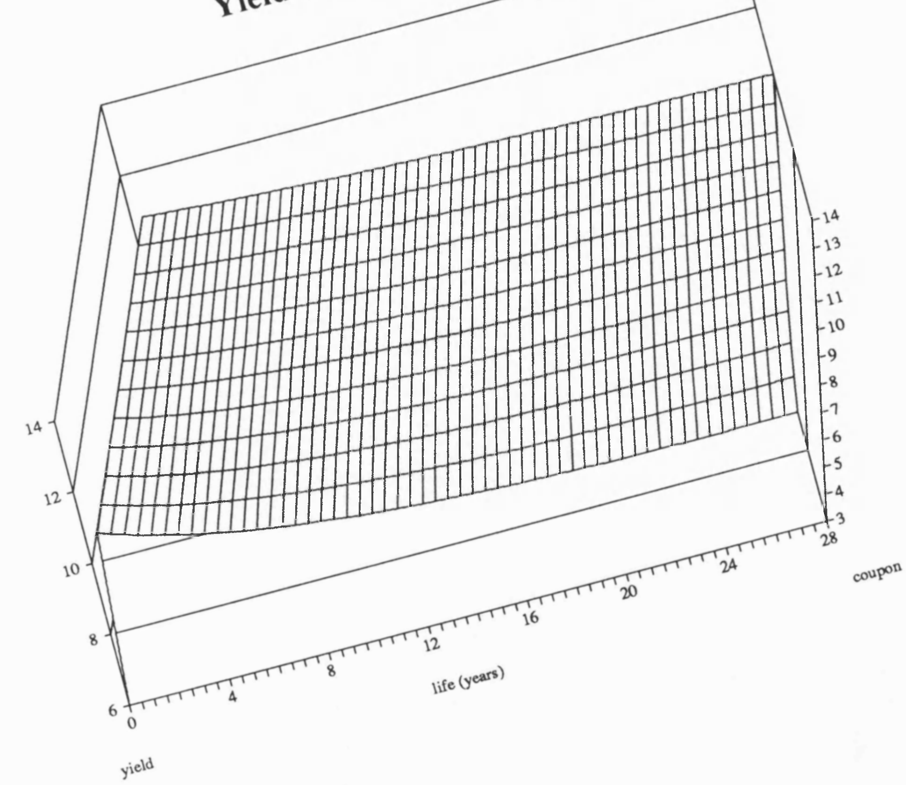

The graph below shows the estimated composite yield surface corresponding to the previous three. In this circumstance, it would appear that the longer maturities are dominated by gross investors, the low coupon gilts - especially at shorter maturities - are dominated by net investors, whereas the short to medium maturities and medium to high coupon gilts are dominated by net net investors. This is also generally true. More details regarding the segmentation are given in the next three chapters.

Yield surface

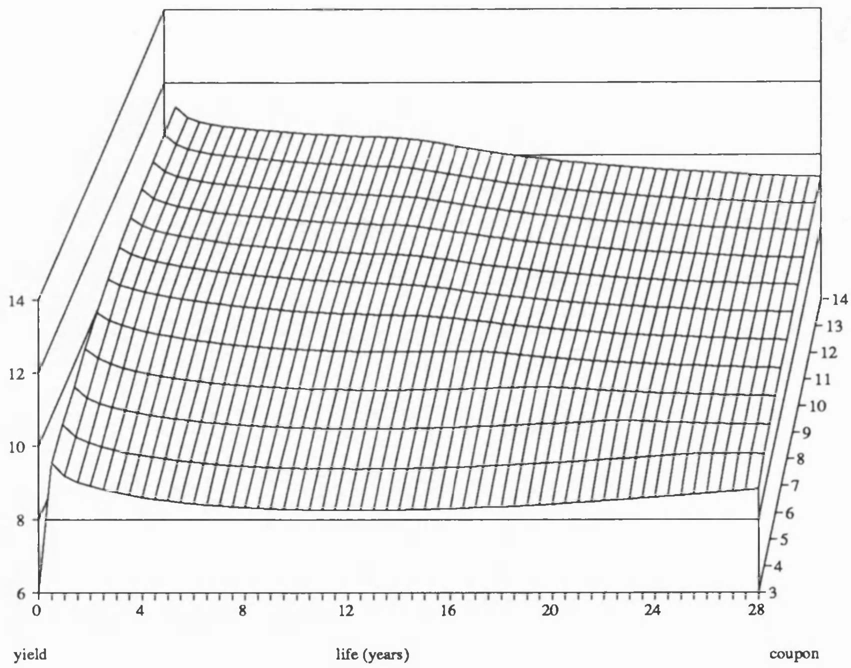

It is worth pointing out at this stage how this situation gives rise to an indeterminacy problem: the discount function for net net investors is poorly determined at long maturities, as is the discount function for gross investors at short maturities, and possibly also the discount functions for net investors at the longest maturities. This issue is discussed further in section 6.2.

IV. UK GILTS: EMPIRICAL RESULTS

4.1 Model estimation

The UK gilt market model was estimated using daily prices over 1987-88. The data was obtained from Morgan Grenfell Government Securities, which at the time was a gilt-edged market maker. As explained in the introduction of chapter III, only dated conventional gilts with no significant option element were included. Issues that were negligibly small (e.g. because they had been almost entirely converted) or exceptionally illiquid due to special circumstances (e.g. 3% Exchequer Gas 1990/95) were not used. Also, recently issued gilts were usually omitted for several days after issue.

The discount functions for each type of investor are specified in the same way as those used in the Bund and Dutch bullet models, but with a different choice of rates $\{r_k\}$ in order to maintain the number of parameters at a manageable level. The rates chosen are:

$r_1 = 3\%$
$r_2 = 6\%$
$r_3 = 12\%$
$r_4 = 24\%$.

The graph below shows the standard error of the gilt model over the two-year period covered.

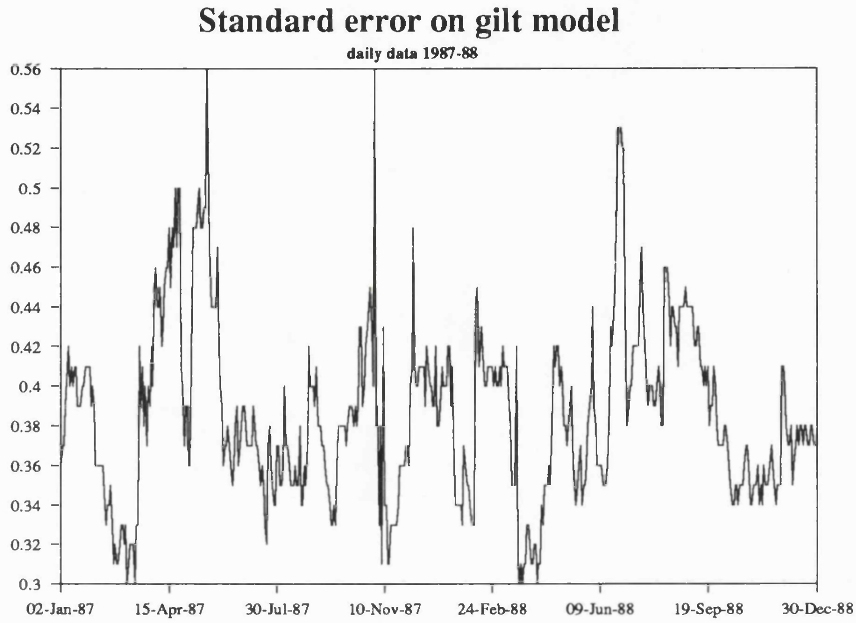

The standard error averages just below 0.40, or around 12/32, over the period. This is a much worse fit than in the case of either the German Bund or the Dutch bullet model. There are several possible reasons for this. Firstly, the greater complexity of modelling a market with taxation effects and, in particular, the weaknesses inherent in calculations of net of tax cash flows. Secondly, a number of poorly traded gilts were included, and the gilt market doesn't benefit from regulatory intervention as in the case of the German Bund market. Thirdly, no special allowance has been made for FOTRA gilts, which tend to trade at a premium. These points are examined in greater detail in later sections of this chapter and in chapter V.

The estimation on data for 1 September 1988, which was fairly average for the period considered, gave the following results:

Table 8. Estimation on 1 September 1988

Gross segment:		
$\beta_1 = -0.137$	$s_1 = 0.062$	$t_1 = -2.2$
$\beta_2 = 0.897$	$s_2 = 0.165$	$t_2 = 5.4$
$\beta_3 = -0.171$	$s_3 = 0.167$	$t_3 = -1.0$
$\beta_4 = 0.412$	$s_4 = 0.064$	$t_4 = 6.4$
Net segment:		
$\beta_5 = -0.173$	$s_5 = 0.135$	$t_5 = -1.3$
$\beta_6 = 1.298$	$s_6 = 0.290$	$t_6 = 4.5$
$\beta_7 = -0.299$	$s_7 = 0.209$	$t_7 = -1.4$
$\beta_8 = 0.174$	$s_8 = 0.055$	$t_8 = 3.1$
Net net segment:		
$\beta_9 = -0.013$	$s_9 = 0.125$	$t_9 = -0.1$
$\beta_{10} = 0.957$	$s_{10} = 0.255$	$t_{10} = 3.7$
$\beta_{11} = 0.012$	$s_{11} = 0.170$	$t_{11} = 0.1$
$\beta_{12} = 0.044$	$s_{12} = 0.040$	$t_{12} = 1.1$
$\bar{R}^2 = 0.998$	SE = 0.44	
observations = 85	degrees of freedom = 76	

Note: model estimated under the constraints:

$$\beta_4 = \left(1 - \sum_{k=1}^{3} \beta_k\right), \quad \beta_8 = \left(1 - \sum_{k=5}^{7} \beta_k\right) \text{ and } \beta_{12} = \left(1 - \sum_{k=9}^{11} \beta_k\right).$$

The low values of the asymptotic t statistics on the coefficients of the net net segment of the market suggest that there could be problems with the model specification in this case. The result is consistent with the previous remarks regarding the timing option of tax on capital gains. This point is discussed in more detail in chapter V.

The graph below overlays coupon yield curves estimated on 1 September 1988 with yields on the same date of gilts that were used in the estimation. The coupon effect is very clear, as is the segmentation. The table that follows provides the key to the graph.

UK gilt yield structure

on 1 September 1988

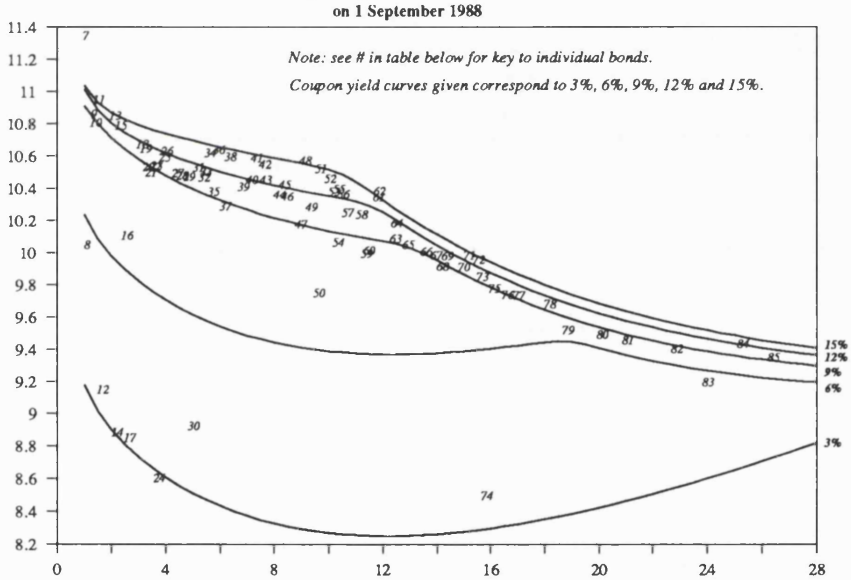

Table 9. UK gilt list on 1 September 1988

#	stock	price	yield
1	9.5 Tr 25-Oct-88	99.59	11.903
2	11.5 Tr 22-Feb-89	99-27	11.705
3	9.5 TC 18-Apr-89	98-24	11.597
4	3 Tr 15-May-89	96-00	9.047
5	10.5 Tr 14-Jun-89	99-07	11.527
6	10 Ex 1-Aug-89	98-24	11.459
7	11 Ex 29-Sep-89	99-22	11.345
8	5 Tr 15-Oct-86/89	94-24	10.051
9	*f* 13 Tr 15-Jan-90	102-18	10.860
10	*f* 11 Ex 12-Feb-90	100-06	10.809
11	12.5 Ex 22-Mar-90	102-07	10.951
12	3 Tr 8-May-90	90-20	9.152
13	10 TC 25-Oct-90	98-12	10.849
14	2.5 Ex 22-Nov-90	87-12	8.885
15	11.75 Tr 10-Jan-91	101-28	10.789
16	*f* 5.75 Fn 5-Apr-87/91	90-10	10.110
17	3 Tr 13-May-91	86-08	8.851
18	11 Ex 25-Oct-91	100-26	10.673
19	8 Tr 10-Dec-91	92-26	10.646
20	*f* 12.75 Tr 22-Jan-92	106-03	10.535
21	10 Tr 21-Feb-92	98-17	10.501
22	*f* 8 Tr 13-Apr-92	92-16	10.536
23	*f* 10.5 TC 7-May-92	99-26	10.542
24	3 Tr 11-Jun-92	82-08	8.605
25	12.25 Ex 25-Aug-92	105-07	10.593
26	13.5 Ex 22-Sep-92	109-08	10.632
27	8.25 Tr 18-Feb-93	92-03	10.498
28	*f* 10 Tr 15-Apr-93	98-07	10.482
29	*f* 12.5 Tr 14-Jul-93	107-16	10.473
30	*f* 6 Fn 15-Sep-93	88-11	8.923
31	*f* 13.75 Tr 23-Nov-93	112-18	10.531
32	8.5 Tr 3-Feb-94	91-30	10.471
33	*f* 14.5 Tr 1-Mar-94	116-12	10.504
34	13.5 Ex 27-Apr-94	111-30	10.618
35	*f* 10 Tr 9-Jun-94	98-10	10.382
36	12.5 Ex 22-Aug-94	108-01	10.637
37	*f* 9 Tr 17-Nov-94	94-04	10.287
38	12 Tr 25-Jan-95	106-10	10.596
39	10.25 Ex 21-Jul-95	99-05	10.412
40	*f* 12.75 Tr 15-Nov-95	111-10	10.454
41	14 Tr 22-Jan-96	117-02	10.590
42	*f* 15.25 Tr 3-May-96	124-06	10.549
43	*f* 13.25 Ex 15-May-96	114-14	10.455
44	10 Cv 15-Nov-96	97-30	10.362
45	*f* 13.25 Tr 22-Jan-97	115-14	10.423
46	10.5 Ex 21-Feb-97	100-24	10.351
47	*f* 8.75 Tr 1-Sep-97	91-22	10.178
48	15 Ex 27-Oct-97	125-14	10.575
49	9.75 Ex 19-Jan-98	96-24	10.283
50	*f* 6.75 Tr 1-May-95/98	81-14	9.751
51	14 Tr 22-May-98/01	120-23	10.525
52	*f* 15.5 Tr 30-Sep-98	130-29	10.463
53	12 Ex 20-Nov-98	109-30	10.380
54	*f* 9.5 Tr 15-Jan-99	96-10	10.067
55	12 Ex 22-Jan-99/02	109-30	10.396
56	12.25 Ex 26-Mar-99	111-30	10.365
57	10.5 Tr 19-May-99	101-16	10.252
58	10.25 Cv 22-Nov-99	99-31	10.240
59	*f* 8.5 Tr 28-Jan-00	89-28	9.996
60	*f* 9 Cv 3-Mar-00	93-04	10.017
61	13 Tr 14-Jul-00	117-25	10.347
62	13.75 Tr 25-Jul-00/03	122-18	10.386
63	10 Tr 26-Feb-01	99-09	10.090
64	11.5 Tr 19-Mar-01/04	109-06	10.185
65	9.75 Cv 10-Aug-01	97-25	10.049
66	10 Cv 11-Apr-02	99-28	10.007
67	9.75 Tr 27-Aug-02	98-06	9.984
68	9 Ex 19-Nov-02	93-00	9.916
69	11.75 Tr 22-Jan-03/07	113-08	9.980
70	10 Tr 8-Sep-03	100-21	9.911
71	12.5 Tr 21-Nov-03/05	119-11	9.985
72	13.5 Tr 26-Mar-04/08	127-23	9.958
73	10 Tr 18-May-04	101-02	9.852
74	3.5 Fn 14-Jul-99/04	56-27	8.499
75	9.5 Cv 25-Oct-04	97-22	9.777
76	9.5 Cv 18-Apr-05	97-31	9.739
77	10.5 Ex 20-Sep-05	106-09	9.734
78	9.75 Cv 15-Nov-06	100-15	9.682
79	*f* 8.5 Tr 16-Jul-07	91-02	9.519
80	*f* 9 Tr 13-Oct-08	95-17	9.492
81	8 Tr 25-Sep-09	86-24	9.459
82	*f* 9 Cv 12-Jul-11	96-04	9.405
83	*f* 5.5 Tr 10-Sep-08/12	64-13	9.197
84	12 Ex 12-Dec-13/17	124-13	9.436
85	*f* 7.75 Tr 26-Jan-12/15	84-10	9.353

Note: *f* denotes FOTRA gilts.

4.2 Effective tax rates and the 1988 budget

An important test of the gilt model is the comparison of the "effective tax rates", i.e. the rates of tax that provide the best fit of the model, with the rates of tax that are known a priori to apply to the different categories of investors. The table below summarizes the results obtained over the period from October 1987 to June 1988, which for reasons described later is particularly interesting.

Table 10. Average standard error of the gilt model
October 1987 - June 1988

profits tax rate	income tax rate										
	20	25	30	35	40	45	50	55	60	65	70
10	0.703	0.58	0.531	0.533	0.569	0.589	0.618	0.65	0.692	0.751	0.815
15	0.686	0.549	0.498	0.514	0.53	0.546	0.568	0.597	0.644	0.704	0.763
20	0.668	0.536	0.481	0.483	0.481	0.489	0.513	0.551	0.602	0.654	0.708
25	0.663	0.523	0.464	0.446	0.438	0.448	0.476	0.52	0.576	0.64	0.69
30	0.666	0.523	0.456	0.427	0.416	0.427	0.457	0.502	0.555	0.619	0.669
35	0.677	0.544	0.451	0.413	0.4	0.411	0.441	0.484	0.54	0.598	0.643
40	0.708	0.574	0.477	0.411	0.398	0.41	0.44	0.479	0.525	0.581	0.631
45	0.732	0.623	0.516	0.446	0.407	0.418	0.447	0.484	0.525	0.573	0.609
50	0.736	0.662	0.596	0.502	0.46	0.437	0.46	0.496	0.538	0.575	0.601

It is apparent that the minimum average standard error over the period considered is achieved for an income tax rate of around 40% and a corporation tax rate of 35-40%. The rate of corporation tax used in the model estimation over the whole of 1987-88 was fixed at 35%, in view of the fact that this was the top rate of corporation tax during that period.

In his 1988/89 Budget speech[1] the Chancellor announced that the top rate of income tax was to be reduced from 60% to 40%. A reduction was generally expected, but not of such a magnitude. Expectation of a cut in the top rate had built up over a period of months as information emerged regarding to the state of public finances and government intentions. A landmark in this process was the previous Autumn Statement.[2] An interesting test of the gilt model is whether this process of expectation formation is reflected in the effective rates of tax. The following table shows very clearly the correct type of response on the part of the model to the tax changes.

Table 11. Average standard error of the gilt model October 1987 - June 1988

trade dates	income tax rates[3]										
	20	25	30	35	40	45	50	55	60	65	70
1987Q4	0.748	0.620	0.476	0.431	0.410	0.410	0.429	0.465	0.534	0.577	0.617
1988Q1	0.654	0.518	0.442	0.405	0.390	0.403	0.438	0.480	0.532	0.590	0.635
1988Q2	0.629	0.493	0.433	0.402	0.398	0.419	0.455	0.505	0.555	0.628	0.676

The table shows how the rate of income tax implied by the model declines from 40-45% in the autumn of 1987 to 35-40% in the spring of 1988. Similar calculations in 1986 give an effective rate of income tax of 45-50%. In view of this, it was decided to use a rate of income tax of 45% in the model estimation up to the budget, and a rate of 35% thereafter.

1. This took place on 15 March 1988.

2. This was delivered on 3 November 1987.

3. The profits tax rate is fixed at 35%.

4.3 Error analysis

The following model, subject to different sets of restrictions, was fitted to the time series of estimated errors:

$$
\begin{cases}
\varepsilon_{t,i} = \mu_i + \beta_i(\varepsilon_{t-1,i} - \mu_i) + \xi_{t,i} \\
E\xi_{t,i} = 0, E\xi_{t_1,i_1}\xi_{t_2,i_2} = \begin{cases} 0 & \text{if} \quad t_1 \neq t_2 \quad \text{or} \quad i_1 \neq i_2 \\ \sigma^2 & \text{otherwise} \end{cases} \\
E\xi_{t,i_1}\mu_{i_2} = 0, E\xi_{t,i_1}\beta_{i_2} = 0, E\xi_{t-1,i_1}\varepsilon_{t,i_2} = 0 \\
E\mu_i = m_\mu, E(\mu_{i_1} - m_\mu)(\mu_{i_2} - m_\mu) = \begin{cases} 0 & \text{if} \quad i_1 \neq i_2 \\ \sigma_\mu^2 & \text{otherwise} \end{cases} \\
E\beta_i = m_\beta, E(\beta_{i_1} - m_\beta)(\beta_{i_2} - m_\beta) = \begin{cases} 0 & \text{if} \quad i_1 \neq i_2 \\ \sigma_\beta^2 & \text{otherwise} \end{cases} \\
E(\mu_{i_1} - m_\mu)(\beta_{i_2} - m_\beta) = 0.
\end{cases}
$$

The estimation, carried out across 91 gilts using daily data for 1987-88, produced:

Table 12. Results of error model estimation

model	$\hat{m}_\mu$	$\hat{\sigma}_\mu$	$\hat{m}_\beta$	$\hat{\sigma}_\beta$	$\hat{\sigma}$	df
A	0†	0†	0†	0†	0.40	40569
B	0†	0†	0.93	0.07	0.10	40478
C	-0.04	0.34	0†	0†	0.22	40478
D	-0.04	0.34	0.84	0.09	0.10	40387
E	0†	0†	1†	0†	0.10	40569
F	0†	0†	0.97†	0†	0.10	40568
G	-0.04	0.34	0.9	0†	0.10	40477

Note: a † indicates a constrained parameter in the estimation.

The main fact that emerges from this table is that the errors appear to be strongly auto-correlated, to the point that they closely resemble a random walk.

As discussed in section 1.4, one would not expect pricing errors to fall outside a band of width comparable to twice the bid-offer spread except in so far as the model failed to capture all the factors affecting investors' judgement of value. If the neglected factors usually changed slowly, errors in price should vary within a band around some slowly changing level (long-term mispricing). In this case arbitrage profits would not be possible even if long-term mispricing tended to revert towards zero, owing to the costs of running a long term position. The following tables illustrate how, by using a simple rule of thumb,[1] it was possible in 1988 to make a return that would just about equal transaction costs over a holding period of 30 working days.

Table 13. Example of portfolio simulation
4 January 1988

stock	nominal amount	price 04-Jan	price change by: 05-Jan	11-Jan	18-Jan	25-Jan	01-Feb	08-Feb	15-Feb
f 6.75 Tr 1-May-95/98	7	83.13	-0.38	-0.81	-0.06	1.19	1.38	0.13	0.44
12 Ex 12-Dec-13/17	9	124.38	-0.31	-1.56	0.44	2.00	1.94	-0.19	0.56
13.75 Tr 25-Jul-00/03	11	126.75	-0.38	-1.56	0.13	1.63	1.50	-0.38	0.06
13 Tr 14-Jul-00	10	121.75	-0.38	-1.50	0.19	1.63	1.56	-0.38	0.06
14 Tr 22-May-98/01	10	125.50	-0.38	-1.56	-0.31	1.19	1.19	-0.75	-0.38
f 8.5 Tr 16-Jul-07	-9	92.06	-0.25	-1.25	0.44	1.44	1.38	-0.31	0.44
f 13.75 Tr 23-Nov-93	-10	118.56	-0.25	-1.19	-0.81	0.44	0.19	-1.19	-1.19
9 Ex 19-Nov-02	-11	95.50	-0.38	-1.13	0.31	1.50	1.31	-0.50	0.19
f 9.5 Tr 15-Jan-99	-11	99.31	-0.38	-1.31	0.00	1.13	1.13	-0.50	0.06
f 14.5 Tr 1-Mar-94	-10	122.72	-0.25	-1.28	-0.59	0.59	0.41	-1.09	-1.16
portfolio gain:			-0.01	-0.03	0.10	0.19	0.25	0.21	0.23

Note: *f* denotes FOTRA gilts.

1. The rule adopted is:

- portfolio duration, running yield and money invested are all constrained to zero
- total transaction nominal is 100
- the 5 most cheap bonds are bought and the 5 most dear bonds are sold in amounts as equal as possible (in a least squares sense) given the above constraints. Cheapness and dearness are defined by (mispricing - average) / standard deviation, with the statistics calculated over the previous six months.

Table 14. Summary portfolio simulation returns daily data 1988

trade date	gain after days: 1	5	10	15	20	25	30
04-Jan-88	-0.01	-0.03	0.10	0.19	0.25	0.21	0.23
05-Jan-88	0.04	0.01	0.05	0.06	0.13	0.09	0.11
06-Jan-88	-0.07	0.00	0.02	0.07	0.18	0.20	0.20
07-Jan-88	-0.01	0.01	0.14	0.19	0.25	0.26	0.33
08-Jan-88	0.04	0.06	0.16	0.21	0.22	0.24	0.30
11-Jan-88	0.03	0.13	0.17	0.22	0.24	0.25	0.39
12-Jan-88	0.04	0.07	0.09	0.21	0.18	0.21	0.29
13-Jan-88	-0.02	0.02	0.07	0.14	0.17	0.16	0.28
14-Jan-88	0.02	0.12	0.16	0.24	0.24	0.32	0.43
15-Jan-88	0.30	0.31	0.38	0.41	0.41	0.43	0.56
18-Jan-88	-0.05	0.02	0.07	0.10	0.12	0.23	0.27
19-Jan-88	0.01	0.05	0.19	0.17	0.16	0.21	0.24
20-Jan-88	0.02	0.04	0.18	0.19	0.17	0.26	0.25
21-Jan-88	0.02	0.03	0.08	0.07	0.11	0.17	0.19
22-Jan-88	0.01	0.04	0.05	0.07	0.16	0.28	0.21
25-Jan-88	-0.04	0.06	0.02	0.07	0.29	0.34	0.30
26-Jan-88	0.02	0.15	0.12	0.16	0.29	0.30	0.36
27-Jan-88	0.16	0.25	0.24	0.28	0.38	0.38	0.41
28-Jan-88	0.07	0.11	0.10	0.19	0.25	0.27	0.36
29-Jan-88	0.00	0.01	0.04	0.21	0.36	0.25	0.49
01-Feb-88	0.06	-0.10	0.00	0.34	0.39	0.36	0.60
02-Feb-88	0.01	-0.05	0.03	0.19	0.20	0.29	0.34
03-Feb-88	-0.03	0.00	0.05	0.20	0.14	0.22	0.21
04-Feb-88	0.01	0.05	0.16	0.21	0.28	0.37	0.20
05-Feb-88	-0.06	0.00	0.12	0.19	0.17	0.29	0.13
08-Feb-88	0.00	0.01	0.11	0.14	0.15	0.22	0.10
09-Feb-88	0.15	0.13	0.18	0.19	0.20	0.28	0.21
10-Feb-88	-0.02	0.00	0.09	0.03	0.08	0.13	0.05
11-Feb-88	0.02	0.04	0.10	0.11	0.21	0.09	0.06
12-Feb-88	-0.03	0.05	0.14	0.08	0.18	0.07	0.04
15-Feb-88	0.02	0.10	0.10	0.09	0.22	0.10	0.13
17-Feb-88	0.02	0.05	0.06	0.09	0.19	0.09	0.06
18-Feb-88	0.04	0.11	0.06	0.11	0.15	0.05	0.16
....	....	....	....	....	....	....	....
16-Nov-88	0.00	0.03	0.03	0.01	0.00	-0.01	-0.02
17-Nov-88	0.00	0.02	0.02	0.04	0.02	0.01	0.01
18-Nov-88	0.01	0.00	-0.02	-0.06	-0.05	-0.02	0.02
21-Nov-88	0.01	-0.02	-0.04	-0.06	-0.07	-0.06	-0.03
22-Nov-88	0.01	-0.02	-0.03	-0.04	-0.04	-0.04	0.02
23-Nov-88	0.00	-0.02	-0.04	-0.04	-0.06	-0.05	0.00
24-Nov-88	0.00	0.01	0.02	0.00	-0.01	-0.01	0.03
25-Nov-88	0.02	0.00	0.02	0.06	0.02	0.02	0.05
28-Nov-88	0.01	-0.01	-0.03	-0.03	-0.03	0.00	-0.01
29-Nov-88	0.00	-0.03	-0.04	-0.08	-0.05	-0.01	-0.04
30-Nov-88	0.00	-0.03	-0.05	-0.08	-0.04	-0.02	-0.04
01-Dec-88	-0.01	-0.03	-0.03	-0.05	-0.03	-0.01	-0.07
02-Dec-88	-0.01	-0.03	-0.05	-0.04	0.00	0.00	-0.07
05-Dec-88	0.02	0.01	-0.03	0.00	0.03	0.01	-0.06
06-Dec-88	-0.01	-0.02	-0.06	-0.02	0.02	-0.01	-0.09
07-Dec-88	0.03	0.00	-0.02	0.02	0.04	0.02	-0.06
08-Dec-88	-0.02	-0.03	-0.05	-0.02	-0.01	-0.06	-0.10
09-Dec-88	0.01	-0.06	-0.02	0.02	0.00	-0.05	-0.08
12-Dec-88	-0.01	-0.05	-0.01	0.02	0.00	-0.06	-0.05
13-Dec-88	-0.02	-0.04	0.00	0.04	0.01	-0.06	-0.06
14-Dec-88	0.01	-0.02	0.02	0.04	0.02	-0.03	-0.02
15-Dec-88	-0.06	-0.02	0.01	0.02	-0.02	-0.05	-0.02
16-Dec-88	0.02	0.04	0.10	0.10	0.07	0.08	0.12
19-Dec-88	0.00	0.05	0.09	0.07	0.04	0.08	0.12
20-Dec-88	0.01	0.06	0.09	0.08	0.05	0.07	0.14
21-Dec-88	-0.01	0.05	0.08	0.12	0.08	0.08	0.09
22-Dec-88	0.01	0.07	0.10	0.12	0.07	0.10	0.09
23-Dec-88	0.04	0.09	0.10	0.10	0.08	0.10	0.08
28-Dec-88	0.02	0.04	0.05	0.02	0.04	0.06	0.03
29-Dec-88	-0.01	0.03	0.05	0.00	0.00	0.04	-0.01
30-Dec-88	0.02	0.03	0.08	0.04	0.04	0.05	-0.03
average:	0.01	0.02	0.04	0.05	0.06	0.07	0.08

Despite the relatively high standard error of the gilt model, it therefore appears that arbitrage profits net of costs might be extremely small. Thus the size of the standard error is likely to be in part a reflection of the inability of the model to capture all the complexities of the gilt market rather than a case of inefficient behaviour. One factor that might help to explain the size of the standard error is the FOTRA effect, which is discussed in the following section.

4.4 The FOTRA effect

It is common belief of gilt market participants that FOTRA stocks should enjoy a premium in view of their privileged tax treatment. A priori, such a premium should be small, since an overseas investor can wash the coupon when necessary on non-FOTRA stocks at an average cost that is almost negligible in view of the fact that the expected holding period is usually quite short. To be more precise, an overseas investor who planned to have a stable holding of gilts would have transaction costs on non-FOTRA gilts of about twice[1] the bid-offer spread each year, which on a 10 year gilt would be equivalent to about 1% in price terms. One would expect the FOTRA effect to be less than that. The graph and table that follow provide some empirical evidence on this subject.

Mean error of FOTRAs vs. duration

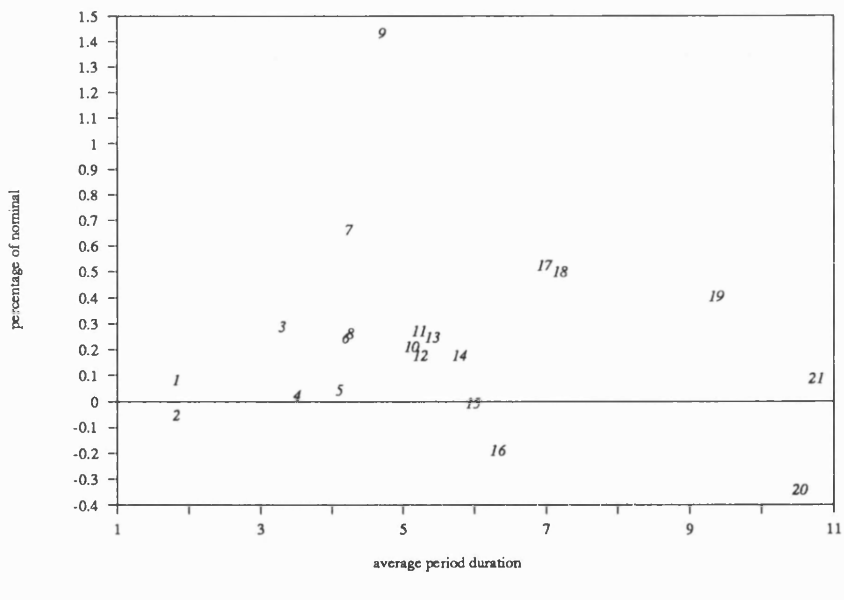

1. Since gilts pay coupons semi-annually, two washes a year would be required.

Table 15. Mispricing of FOTRA gilts daily data 1987-88					
#	stock	average duration	mispricing statistics mean	std dev	std var
1	13 Tr 15-Jan-90	1.82	0.08	0.14	0.16
2	11 Ex 12-Feb-90	1.83	-0.05	0.14	0.15
3	12.75 Tr 22-Jan-92	3.30	0.29	0.11	0.31
4	10.5 TC 7-May-92	3.51	0.03	0.11	0.11
5	10 Tr 15-Apr-93	4.10	0.05	0.16	0.16
6	12.5 Tr 14-Jul-93	4.19	0.25	0.12	0.27
7	14.5 Tr 1-Mar-94	4.23	0.67	0.13	0.68
8	13.75 Tr 23-Nov-93	4.25	0.26	0.15	0.30
9	6 Fn 15-Sep-93	4.71	1.43	0.58	1.55
10	9 Tr 17-Nov-94	5.12	0.21	0.25	0.33
11	15.25 Tr 3-May-96	5.22	0.27	0.15	0.31
12	12.75 Tr 15-Nov-95	5.24	0.18	0.13	0.22
13	13.25 Ex 15-May-96	5.41	0.25	0.11	0.27
14	13.25 Tr 22-Jan-97	5.80	0.18	0.13	0.22
15	15.5 Tr 30-Sep-98	5.99	-0.01	0.29	0.29
16	8.75 Tr 1-Sep-97	6.33	-0.19	0.27	0.33
17	9.5 Tr 15-Jan-99	6.97	0.53	0.14	0.54
18	9 Cv 3-Mar-00	7.19	0.50	0.15	0.52
19	8.5 Tr 16-Jul-07	9.37	0.41	0.15	0.44
20	7.75 Tr 26-Jan-12/15	10.52	-0.34	0.28	0.44
21	5.5 Tr 10-Sep-08/12	10.76	0.09	0.55	0.56
	FOTRA average	5.52‡	0.24‡		
			0.43†	0.24†	0.49†
	non-FOTRA average	5.38‡	-0.10‡		
			0.25†	0.16†	0.30†
	total average	5.42‡	0.01‡		
			0.32†	0.19†	0.37†

Note: ‡ = arithmetical mean, † = root mean square.

There are a number of facts that emerge from the analysis above. Firstly, the average mispricing of FOTRA relative to non-FOTRA gilts is around 1/3, i.e. over 3 times the average bid-offer spread, but is well within the levels one would expect a priori. Secondly, the effect doesn't appear to be increasing in duration, as one might expect;[1] this is probably due to the fact that foreign investors buy selectively, and that therefore the FOTRA effect is a function of stock and of foreign

1. The regression of average mispricing against duration produces a negative (and insignificant) coefficient and a highly significant constant

investor interest. Thirdly, a crude calculation[1] suggests that eliminating the FOTRA effect would lead to a standard error comparable to that recorded on Bund, so that the relatively poor performance of the gilt model is probably due to the FOTRA effect; unfortunately, it is almost impossible within a discount function framework to take the FOTRA effect into account in a way that isn't either highly contrived or much too parameterized to be feasible, and this has not therefore been attempted.[2] Fourthly, the FOTRA effect appears to be relatively volatile, since the standard deviation of FOTRA gilt errors is 50% higher than that of non-FOTRA gilts.

1. Bearing in mind that there are 21 FOTRA gilts with average standard error of 0.49, and 45 non-FOTRA gilts with average standard 0.30, by solving for x the system:

$$\begin{cases} 21 \cdot 0.49^2 = 21 \cdot x^2 + f \\ 45 \cdot 0.30^2 = 45 \cdot x^2 + f \end{cases}$$

one gets a guess at the net of FOTRA standard error of 0.21.

2. It is worth pointing out that the FOTRA effect, while large compared to the model standard error, is small when compared to the standard errors of other comparable models in the literature.

4.5 Heteroscedasticity & the bid-offer spread

It has been argued in the literature, on somewhat flimsy grounds, that errors are homoscedastic in yield terms. As explained in section 1.4, a rather more convincing argument is that the error variances might depend on the bid-offer spread. A statistical investigation of this fact is difficult because reliable data on bid-offer spreads is not available, and because the spreads vary over time and from market maker to market maker. The data reported in the table on the next page are only broadly indicative.

Error standard deviation vs duration

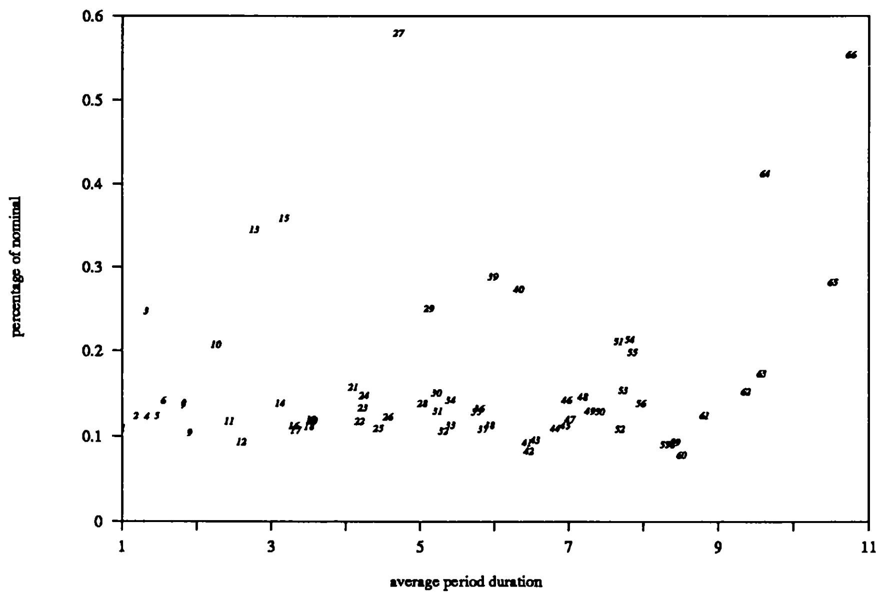

Table 16. Mispricing and bid-offer spread daily data 1987-88

#	stock	average duration	bid-offer spread	standard deviation
1	11.5 Tr 22-Feb-89	1.014	0.03125	0.110
2	9.5 TC 18-Apr-89	1.183	0.06250	0.124
3	3 Tr 15-May-89	1.328	0.18750	0.248
4	10.5 Tr 14-Jun-89	1.330	0.03125	0.123
5	10 Ex 1-Aug-89	1.466	0.03125	0.124
6	11 Ex 29-Sep-89	1.547	0.03125	0.142
7	*f* 13 Tr 15-Jan-90	1.818	0.06250	0.137
8	*f* 11 Ex 12-Feb-90	1.826	0.06250	0.139
9	12.5 Ex 22-Mar-90	1.903	0.06250	0.104
10	3 Tr 8-May-90	2.253	0.25000	0.208
11	10 TC 25-Oct-90	2.435	0.06250	0.118
12	11.75 Tr 10-Jan-91	2.602	0.06250	0.094
13	2.5 Ex 22-Nov-90	2.775	0.37500	0.346
14	11 Ex 25-Oct-91	3.122	0.06250	0.139
15	3 Tr 13-May-91	3.174	0.37500	0.359
16	*f* 12.75 Tr 22-Jan-92	3.302	0.06250	0.113
17	10 Tr 21-Feb-92	3.326	0.06250	0.108
18	*f* 10.5 TC 7-May-92	3.511	0.06250	0.112
19	12.25 Ex 25-Aug-92	3.534	0.06250	0.120
20	13.5 Ex 22-Sep-92	3.550	0.06250	0.119
21	*f* 10 Tr 15-Apr-93	4.103	0.06250	0.158
22	*f* 12.5 Tr 14-Jul-93	4.194	0.06250	0.118
23	*f* 14.5 Tr 1-Mar-94	4.233	0.12500	0.134
24	*f* 13.75 Tr 23-Nov-93	4.252	0.06250	0.148
25	13.5 Ex 27-Apr-94	4.443	0.06250	0.110
26	12.5 Ex 22-Aug-94	4.573	0.06250	0.123
27	*f* 6 Fn 15-Sep-93	4.706	0.37500	0.580
28	12 Tr 25-Jan-95	5.038	0.06250	0.139
29	*f* 9 Tr 17-Nov-94	5.120	0.06250	0.251
30	*f* 15.25 Tr 3-May-96	5.223	0.06250	0.151
31	*f* 12.75 Tr 15-Nov-95	5.243	0.06250	0.130
32	14 Tr 22-Jan-96	5.325	0.09375	0.107
33	*f* 13.25 Ex 15-May-96	5.413	0.09375	0.114
34	10.25 Ex 21-Jul-95	5.421	0.06250	0.143
35	15 Ex 27-Oct-97	5.762	0.09375	0.129
36	*f* 13.25 Tr 22-Jan-97	5.795	0.09375	0.133
37	10.5 Ex 21-Feb-97	5.850	0.09375	0.109
38	10 Cv 15-Nov-96	5.940	0.09375	0.114
39	*f* 15.5 Tr 30-Sep-98	5.986	0.09375	0.289
40	*f* 8.75 Tr 1-Sep-97	6.329	0.09375	0.274
41	12.25 Ex 26-Mar-99	6.447	0.09375	0.093
42	12 Ex 20-Nov-98	6.464	0.09375	0.083
43	9.75 Ex 19-Jan-98	6.557	0.09375	0.096
44	10.5 Tr 19-May-99	6.821	0.09375	0.110
45	13 Tr 14-Jul-00	6.963	0.09375	0.113
46	*f* 9.5 Tr 15-Jan-99	6.974	0.09375	0.143
47	10.25 Cv 22-Nov-99	7.021	0.09375	0.120
48	*f* 9 Cv 3-Mar-00	7.191	0.09375	0.147
49	10 Tr 26-Feb-01	7.284	0.12500	0.130
50	9.75 Cv 10-Aug-01	7.417	0.12500	0.129
51	13.5 Tr 26-Mar-04/08	7.670	0.09375	0.212
52	10 Cv 11-Apr-02	7.682	0.12500	0.109
53	9.75 Tr 27-Aug-02	7.730	0.09375	0.154
54	12.5 Tr 21-Nov-03/05	7.810	0.09375	0.214
55	11.75 Tr 22-Jan-03/07	7.848	0.09375	0.199
56	10 Tr 8-Sep-03	7.960	0.12500	0.139
57	10 Tr 18-May-04	8.277	0.15625	0.091
58	10.5 Ex 20-Sep-05	8.351	0.15625	0.091
59	9.5 Cv 25-Oct-04	8.416	0.15625	0.094
60	9.5 Cv 18-Apr-05	8.503	0.15625	0.079
61	9.75 Cv 15-Nov-06	8.819	0.18750	0.125
62	*f* 8.5 Tr 16-Jul-07	9.369	0.15625	0.153
63	8 Tr 25-Sep-09	9.578	0.15625	0.174
64	12 Ex 12-Dec-13/17	9.618	0.18750	0.413
65	*f* 7.75 Tr 26-Jan-12/15	10.525	0.25000	0.282
66	*f* 5.5 Tr 10-Sep-08/12	10.756	0.18750	0.554

Note: *f* denotes FOTRA gilts.

Regressing the standard error on a constant, duration and bid-offer spread one gets the following results:

Table 17. Regression results			
model	constant	duration	bid-offer spread
A	0.07*(2.9)*	-0.00*(-0.0)*	0.87*(6.9)*
B	0.13*(4.5)*	0.01*(1.4)*	0†

Note: t-ratios in round brackets, † = constrained parameter.

This data supports the argument of section 1.4, with the error standard deviation depending on a constant and on the bid-offer spread of the gilt in question, and not on duration.

It is worth pointing out that accounting for the standard deviation of errors doesn't make much difference to the results. The following table compares estimates of the gilt model under homoscedasticity and heteroscedasticity according to the historical standard deviations. The test statistics show that the heteroscedastic estimates are not significantly different form the homoscedastic ones.

Table 18. Gilt model estimation with heteroscedasticity on 1 September 1988		
homoscedastic estimate	**heteroscedastic estimate**	**t-ratio**
Gross segment:		
$\beta_1 = -0.137$	$\beta_1 = -0.149$	$t_1 = -0.2$
$\beta_2 = 0.897$	$\beta_2 = 0.933$	$t_2 = 0.2$
$\beta_3 = -0.171$	$\beta_3 = -0.216$	$t_3 = -0.3$
$\beta_4 = 0.412$	$\beta_4 = 0.431$	$t_4 = 0.3$
Net segment:		
$\beta_5 = -0.173$	$\beta_5 = -0.053$	$t_5 = 0.9$
$\beta_6 = 1.298$	$\beta_6 = 1.049$	$t_6 = -0.9$
$\beta_7 = -0.299$	$\beta_7 = -0.141$	$t_7 = 0.8$
$\beta_8 = 0.174$	$\beta_8 = 0.146$	$t_8 = -0.5$
Net net segment:		
$\beta_9 = -0.013$	$\beta_9 = -0.029$	$t_9 = -0.1$
$\beta_{10} = 0.957$	$\beta_{10} = 0.985$	$t_{10} = 0.1$
$\beta_{11} = 0.012$	$\beta_{11} = -0.000$	$t_{11} = -0.1$
$\beta_{12} = 0.044$	$\beta_{12} = 0.045$	$t_{12} = 0.0$
Overall:		$F_{9,76} = 0.7$

4.6 Conclusions

Statistical and anecdotal evidence that the gilt market is characterized by taxation effects and segmentation appears to be supported by econometric evidence. Chosing a representative investor for each of three major grouping, and assuming segmented net of tax discount functions, provides a sensible representation of these features. The main theoretical weakness of this methodology is the buy-and-hold assumption used for calculating net of tax cash flows. This issue is discussed in the next chapter.

At a more specific level, the FOTRA effect is a major problem with the gilt model, but is of little general interest and not really amenable to modelling, so no attempt has been made formally to account for it. If the FOTRA effect were eliminated, it would appear that the standard error on gilts might be about twice the bid-offer spread, which is the average trading cost of a round-trip in and out of a gilt.

On a more general note, the theoretical advantages of modelling bond markets in the price space are confirmed by the relationship that emerges between error standard deviations and bid-offer spread, which allows rejection of the homoscedasticity in yield theory. Accounting for heteroscedasticity does not, however, produce significantly different estimates.

V. UK GILTS: THE TAX-TIMING OPTION

5.1 Introduction

As already discussed in previous sections, one of the most important theoretical difficulties with net of tax discount function models of bond markets relates to the timing of tax payments on capital gains. In the case of the gilt market, capital gains are usually subject only to corporation tax. The tax-timing option can therefore apply only to the net net segment of the market. Within this segment the bond holder who has not already managed to avoid all tax liability[1] will, neglecting transaction costs,[2] optimally realize a capital loss immediately,[3] and has an incentive to lock in any capital gain.

If a tax-timing option is allowed for, the precise shape of the price surface within the net net segment depends on much stronger assumptions regarding the underlying economy in general and investor behaviour in particular than is required by a buy-and-hold discount function model. In view of this, no attempt is made here to provide a formal model of the tax-timing option.[4] Instead, empirical evidence is sought regarding certain general implications of the tax-timing option approach.

---

1. A discussion of strategies whereby capital losses may be used to defer taxation indefinitely is contained in Constantinides and Scholes [17]. There are a number of legal and practical constraints to the use of these strategies in the UK which ensure that investors are seldom able to escape being taxed altogether.

2. As explained in a later section, the value of the tax-timing option under plausible assumptions is substantial, and accounting for transaction costs would not significantly modify the conclusions of this chapter.

3. This type of consideration leads to the practice known as "bed and breakfasting", i.e. the sale and immediate re-purchase of a security. Dealers usually quote half or two thirds of the bid-offer spread on the round trip for this type of transaction.

4. Such a model is described in Constantinides and Ingersoll [16]. See section 1.2.3 for further references.

One robust implication of the tax-timing option is that the price surface within the net net segment must be convex with respect to the coupon, rather than linear.[1] In fact, given three gilts with the same maturity but different coupons, the cash flows corresponding to the one with intermediate coupon can be exactly replicated by a portfolio of the other two subject to the additional constraint that the tax-timing options be exercised simultaneously, i.e. that the portfolio be sold (or redeemed) in one block. If such a constraint existed, the surface would be linear, whereas the absence of such a constraint implies stochastic dominance of the intermediate stock by the portfolio, and so the surface must be convex.[2] Unfortunately, this is analogous to the effect of segmentation, and may not always be distinguishable empirically from the latter.

Another fairly robust implication of the tax-timing option is that the slope of the price surface with respect to coupon must be steeper than would otherwise be predicted. This is due to the fact that, broadly speaking, the value of the option should increase with the likelihood of capital losses.

---

1. This fact is also exploited in Litzenberger and Rolfo [32] who, however, draw almost opposite conclusions by examining triplets of government bonds with same maturity date and different coupons.

2. This is analogous to the well known result that a portfolio of options on individual assets is worth more than an option on the portfolio of assets.

5.2 Segmentation

The table below provides a summary description of the pattern of gilt market segmentation in 1988. The calculations are based on theoretical bonds priced according to the estimated segmented discount function model. At any given time a bond is classified as belonging to a particular segment if the corresponding investor type provides the highest price. The table is then compiled from the frequency distribution obtained by recalculating the segmentation daily.

Table 19. Pattern of segmentation daily data 1988

	coupon						
	2-4	4-6	6-8	8-10	10-12	12-14	14-16
life							
0-4	*n*	*n*	*n*	*m*	*m*	*m*	*m*
4-8	*n*	*n*	*n*	*m*	*m*	*m*	*m*
8-12	*n*	*n*	*n*	*m*	*m*	*m*	
12-16	*n*	*n*			*g*	*g*	*g*
16-20	*n*	*n*		*g*	*g*	*g*	*g*
20-24	*n*		*g*	*g*	*g*	*g*	*g*
24-28			*g*	*g*	*g*	*g*	*g*

Note: *n* = 75% of the time net, *g* = 75% of the time gross, *m* = 75% of the time net net.

A clearly defined pattern emerges from this table, with the short maturity medium to high coupon sector obviously dominated by net net investors. The next section attempts to provide some indication as to the extent of the effects one might expect to occur in the net net segment due to the tax-timing option.

5.3 The option value

The following table contains the results of a Monte-Carlo simulation aimed at providing a crude indication[1] as to the order of magnitude of the tax-timing option. The figures given are the average net present value[2] of the difference between the capital gains tax payments under the buy-and-hold assumption, and the simulated capital gains tax payments assuming realization of available losses once a year.[3]

Simulated future prices used to obtain these results are generated by the model:[4]

$$\begin{cases} p_t = p_{t-1} + (100 - p_{t-1})/(T - t) + [\partial q(y_t, b_t)/\partial y_t]\, \varepsilon_t \\ E\varepsilon_t = 0, E\varepsilon_t^2 = 1 \end{cases}$$

where t is now, and T is time of maturity, both measured in years.[5]

In the table, areas of the yield surface that did not contain gilts belonging to the net net segment are denoted by n/a. The values obtained are roughly consistent with those reported, under somewhat different conditions, in Constantinides and Ingersoll [16].

1. A deliberate attempt is made to obtain a conservative estimate, so as to bias the results against the conclusions.

2. The discounting rate used is 10%, broadly in line with the gross yields prevailing at the time.

3. This ensures that the benefits of the strategy are not swamped by trading costs. Clearly, one may do better by trading more or less frequently. See the discussion in Constantinides and Ingersoll [16].

4. This is inconsistent with a discount function model, but is probably good enough for the purpose.

5. The assumed standard deviation of yield shocks of 1% per annum is relatively low by historical comparison.

Table 20. Value of the tax-timing option simulation on average 1988 data

	coupon						
	2 ⊣ 4	4 ⊣ 6	6 ⊣ 8	8 ⊣ 10	10 ⊣ 12	12 ⊣ 14	14 ⊣ 16
life							
0 ⊣ 4	n/a	n/a	n/a	0.01	0.03	0.09	n/a
4 ⊣ 8	n/a	n/a	n/a	0.25	0.48	0.87	1.31
8 ⊣ 12	n/a	n/a	n/a	0.82	1.21	1.99	2.76
12 ⊣ 16	n/a	n/a	n/a	n/a	n/a	n/a	n/a
16 ⊣ 20	n/a	n/a	n/a	n/a	n/a	n/a	n/a
20 ⊣ 24	n/a	n/a	n/a	n/a	n/a	n/a	n/a
24 ⊣ 28	n/a	n/a	n/a	n/a	n/a	n/a	n/a
0 ⊣ 12	n/a	n/a	n/a	0.28	0.42	0.84	1.62

For comparative purposes, the next table gives the value of the option to amortize linearly the capital loss, which corresponds to the optimal strategy for the simulation model if yield volatility were set to zero. As one would expect, the intrinsic value increases with coupon and maturity, whereas the time value decreases.

Table 21. Intrinsic value of the tax-timing option based on average 1988 data

	coupon						
	2 ⊣ 4	4 ⊣ 6	6 ⊣ 8	8 ⊣ 10	10 ⊣ 12	12 ⊣ 14	14 ⊣ 16
life							
0 ⊣ 4	n/a	n/a	n/a	0.00	0.03	0.09	n/a
4 ⊣ 8	n/a	n/a	n/a	0.00	0.33	0.82	1.29
8 ⊣ 12	n/a	n/a	n/a	0.00	0.63	1.63	2.57
12 ⊣ 16	n/a	n/a	n/a	n/a	n/a	n/a	n/a
16 ⊣ 20	n/a	n/a	n/a	n/a	n/a	n/a	n/a
20 ⊣ 24	n/a	n/a	n/a	n/a	n/a	n/a	n/a
24 ⊣ 28	n/a	n/a	n/a	n/a	n/a	n/a	n/a
0 ⊣ 12	n/a	n/a	n/a	0.00	0.25	0.80	1.59

5.4 Empirical results

The following table makes it possible broadly to match the pattern of segmentation and predicted tax-timing option effect with that of average model mispricings.

Table 22. Pattern of mean errors
daily data 1988

	coupon						
	2 ⊣ 4	**4 ⊣ 6**	**6 ⊣ 8**	**8 ⊣ 10**	**10 ⊣ 12**	**12 ⊣ 14**	**14 ⊣ 16**
life							
0 ⊣ 4	-0.04*(0.15)*	-0.16*(0.25)*	-0.68*(0.25)*	-0.11*(0.15)*	-0.09*(0.13)*	0.08*(0.17)*	n/a
4 ⊣ 8	-0.89*(0.35)*	1.53*(0.35)*	-0.43*(0.35)*	-0.09*(0.15)*	-0.23*(0.22)*	0.09*(0.12)*	0.56*(0.25)*
8 ⊣ 12	n/a	n/a	-0.65*(0.35)*	0.06*(0.14)*	-0.24*(0.16)*	-0.19*(0.15)*	-0.04*(0.22)*
12 ⊣ 16	0.00*(0.35)*	n/a	n/a	0.13*(0.13)*	0.05*(0.25)*	-0.29*(0.18)*	n/a
16 ⊣ 20	-0.47*(0.35)*	n/a	n/a	0.07*(0.16)*	0.09*(0.35)*	0.17*(0.35)*	n/a
20 ⊣ 24	n/a	0.24*(0.35)*	0.15*(0.35)*	0.16*(0.25)*	n/a	n/a	n/a
24 ⊣ 28	n/a	0.35*(0.35)*	-0.48*(0.35)*	n/a	-0.61*(0.35)*	n/a	n/a
0 ⊣ 12				**-0.04*(0.09)***	**-0.16*(0.09)***	**0.01*(0.09)***	**0.21*(0.17)***

Note: the estimate of the standard deviation of the mean for each class is given in round brackets. The estimates are based on the assumption that the errors can be described by the model:

$$\begin{cases} \varepsilon_{t,i} = \mu_i + \xi_{t,i} \\ E\xi_{t,i} = 0, E\xi_{t_1,i_1}\xi_{t_2,i_2} = \begin{cases} 0 & \text{if } t_1 \neq t_2 \text{ or } i_1 \neq i_2 \\ \sigma^2 & \text{otherwise} \end{cases} \\ E\xi_{t,i_1}\mu_{i_2} = 0 \\ E\mu_i = 0, E\mu_{i_1}\mu_{i_2} = \begin{cases} 0 & \text{if } i_1 \neq i_2 \\ \sigma_\mu^2 & \text{otherwise.} \end{cases} \end{cases}$$

If in this case the tax-timing option approach provided a valid description of investor behaviour one would expect the mean error to be positive in the region with 14 ⊣ 16% coupon and 0 ⊣ 12 year life.

This appears to be so, but the error is small compared to the estimated standard deviation, to the bid-offer spread,[1] and to the level a tax-timing option model would be likely to predict. One would also expect the difference between the mean errors in the region with 8 ⊣ 10% or 14 ⊣ 16% coupon and 0 ⊣ 12 year life, and in the region with 10 ⊣ 14% coupon and 0 ⊣ 12 year life to be positive. But the difference is only 0.09, which is about half what can be accommodated by the bid-offer spread or two thirds of the estimated standard error. Thus, whilst there is some indication of a tax-timing option effect, this is not overall significant and appears to be comparable in size to the bid-offer spread.

1. The average bid-offer spread in that region is almost 0.1.

5.5 Conclusions

The tax-timing option theory has some robust implications. In particular, the value of the option should be considerable, and such that it would be detected statistically if it were reflected in the market. The gilt market, however, appears not to recognize any significant value to the option, although there are signs of a small mispricing of gilts in the direction predicted by the option.

There are a number of reasons why investors might be reluctant to pay a premium for the option:

(i) many investors are reluctant to produce low profits figures, even though this may be tax efficient

(ii) too open tax avoidance could be disallowed by the tax authority

(iii) many investors frequently liquidate holdings independently of tax considerations so as to re-balance their portfolio or to engage in switches aimed at enhancing profits.

Nevertheless these results are somewhat surprising and one cannot help thinking that some degree of inefficiency might be present.

VI. UK GILTS: TERM STRUCTURE ANALYSIS

6.1 Introduction

A discount function model provides a means of describing a bond market equilibrium. The interaction of supply and demand factors should, therefore, be visible in the shape and behaviour of model term structures. It is not easy to provide evidence regarding these interactions because their nature is often long term, and unobservable expectations clearly play a major role. In the case of the UK gilt model, the three net of tax discount functions provide more insight into the determinants of the equilibrium than would usually be the case in a non-segmented market. This chapter provides some evidence of supply and demand interplay, and illustrates how it may be possible to exploit the simultaneous existence of several term structures.

The approach to discount function models adopted in this study is based on an asymmetry between supply and demand. The demand is modelled insofar as required to ensure the existence of a discount function. The supply of stock is taken to be exogenous, and thus negative positions are not allowed. The applicability of such models is, therefore, constrained to certain types of market: for example, they would not usually be appropriate for a corporate bond market. Even if the government's behaviour is exogenous to the model, it clearly is a major determinant of market prices. Apart from long term effects, one would expect the announcement of a new issue to cause localized "ripples" in the market. Empirical evidence regarding this special aspect is discussed in the next chapter.

6.2 Net term structures

The typical shapes of the estimated net of tax term structures[1] generated by the UK gilt model are summarized by the following two graphs.

As mentioned in section 3.8 these term structures are not well defined at certain maturities. This occurs because the (theoretical) price of a stock is a function only of coupon, and of the annuity and zero coupon prices at the same maturity and in the same segment. A set of bond prices can not therefore uniquely determine the price surface (or, equivalently, the term structure) at maturities longer than that of the bond of longest maturity, nor shorter than that of the bond of shortest maturity therein.[2]

To highlight this fact the portions of the term structure curves that do not encompass any actual bonds that fall within the relevant segment are denoted by thinner lines in the graphs. One might take the view that the degree of indeterminacy of the short maturity portion of the gross segment's term structure is in a sense less than that of the long maturity portion of the net and net net segments' term structures.

1. Each term structure is defined in the usual way as the set of zero-coupon (net) yields equivalent to each net discount function.

2. Given the prices of two bonds with different coupons and same maturity t, the price of any bond of maturity t and only the annuity price at maturity $t-1$ can be calculated.

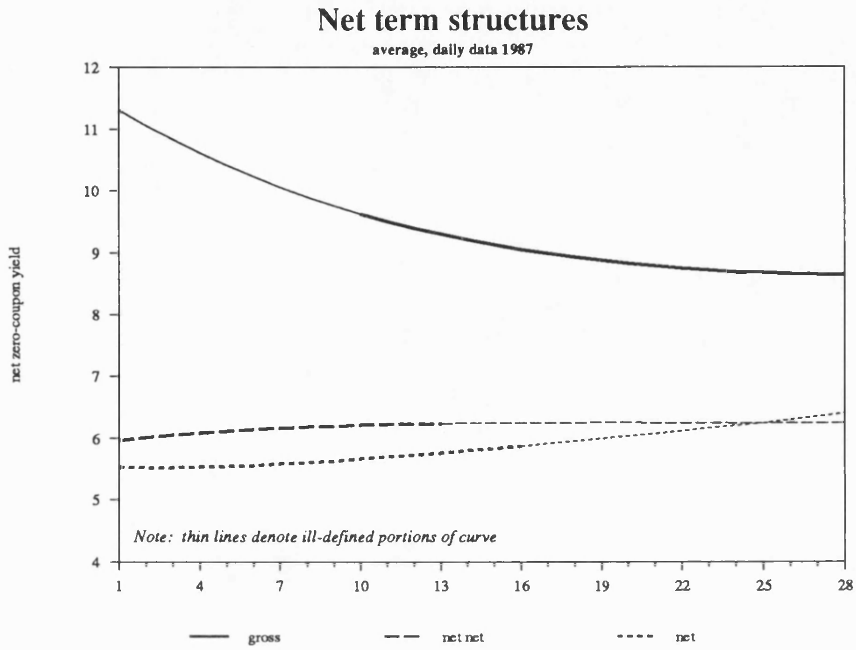

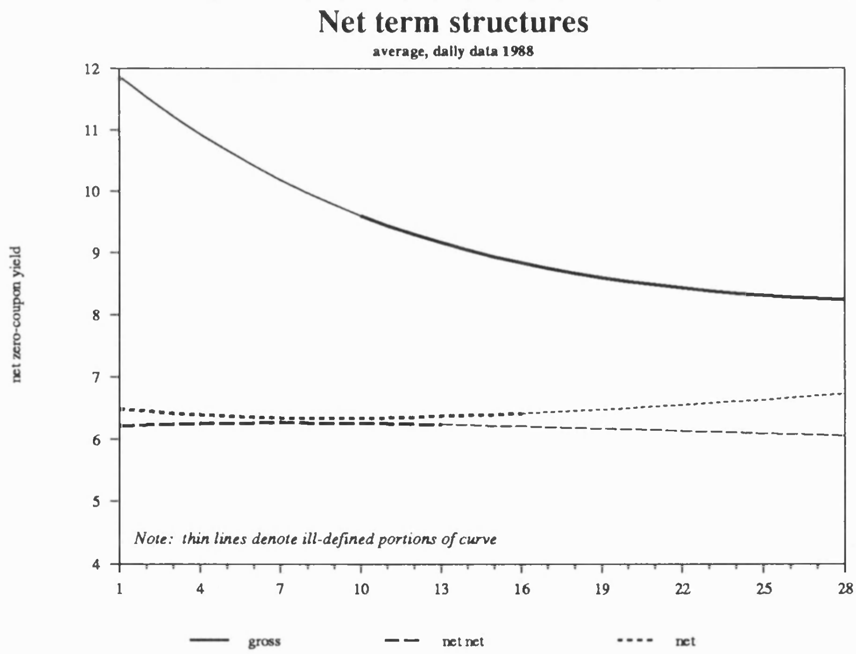

It is apparent from these graphs how relatively complicated price (or yield) surfaces can be generated by combining very simple and plausible net term structures in a model that allows for segmentation.

A number of interesting features emerge from the graphs. It appears for example that, relative to net and net net investors, gross investors have a preference for long maturities. Bearing in mind that most gross investors are pension funds, which have long term liabilities to match and therefore minimise risk by holding bonds of long maturities, this situation is consistent with a situation in which gross investors' appetite for long maturities is not fully met with supply of long gilts but also, partly, with gilts of intermediate maturity.

In these conditions, a reduction in the availability of long gilts should be expected, other things being equal, to produce a more negative slope in the segment's term structure. A comparison of the two graphs shows that the gross segment's term structure underwent precisely such a movement from 1987 to 1988. The following tables tend to confirm this explanation.

Table 23. Nominal gilt market break-down in 1987-88[1]

£ billion *(percent)*

	start 1987	end 1988	change
gross	29*(25)*	22*(18)*	-7*(-25)*
net net	66*(57)*	75*(63)*	9*(14)*
net	9*(8)*	9*(8)*	0*(0)*
other	11*(10)*	13*(11)*	2*(20)*
total	115	120	5*(4)*

1. Only conventional dated gilts are included in this break-down.

Table 24. Nominal gilt holdings break-down in 1987-88			
	£ billion *(percent)*		
	31 March 87	31 March 88	change
pension funds	29*(21)*	26*(19)*	-3*(-11)*
overseas investors	14*(11)*	18*(13)*	3*(22)*
personal sector	15*(11)*	13*(9)*	-1*(-9)*
life assurance	32*(23)*	31*(22)*	-1*(-3)*
monetary sector	9*(6)*	5*(3)*	-4*(-48)*
building societies	9*(6)*	7*(3)*	-1*(-13)*
other[1]	24*(17)*	34*(25)*	11*(45)*
total	137	139	3*(2)*

Sources: CSO, Bank of England, my estimates.

The first table gives a break-down of the gilt market according to the pattern of segmentation described in section 5.2. In interpreting this table it should be born in mind that a stock belongs to a particular segment of the market because of the dominant investor, i.e. the investor who usually buys that type of stock, and not because the stock is held only by a given type of investor. It is immediately apparent that during the period in question there was a substantial fall in the supply of stocks to the gross segment and an equally substantial rise in the supply of stocks to the net net segment.

The second table is not quite comparable to the previous one because it relates to all gilts, rather than only conventional dated gilts, and is based on data as on 31 March. However the picture that emerges is

1. Includes official holdings. These consisted almost exclusively of short and medium maturity gilts.

that of pension funds and, to a lesser extent, life funds being squeezed, with overseas investors increasing their share.[1] Unfortunately, the size of the residual item makes it difficult to interpret the figures further.[2]

Another major feature that emerges from the term structure graphs is the relative rise in the net segment's term structure. The relationship of the latter to that of the net net segment is particularly interesting, and emerges more clearly from the following graphs, which also illustrate further the point made about the gross segment.

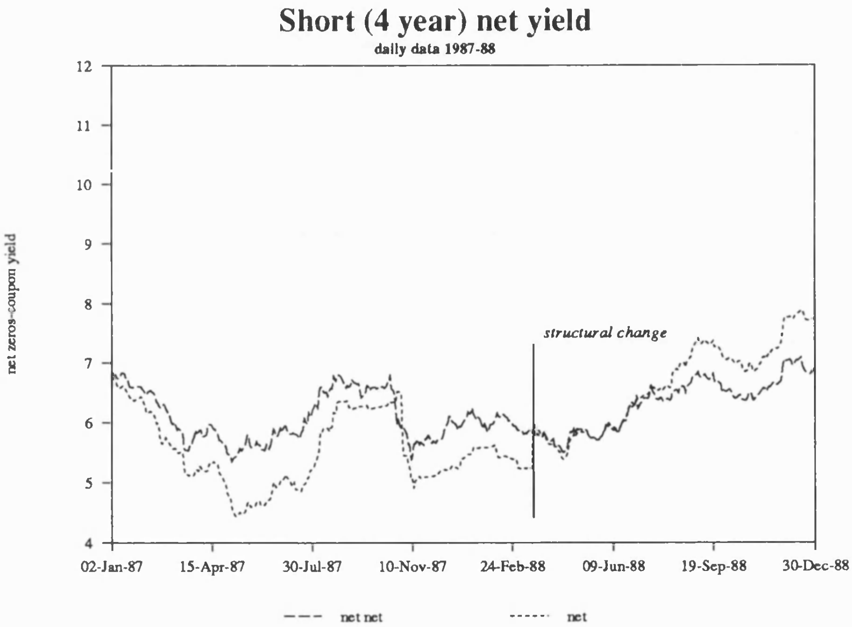

---

1. Overseas investors typically hold mainly short and medium maturity gilts. See section 3.2.1.

2. This item consists almost entirely of short and medium maturity gilts.

Medium (11 year) net yield

daily data 1987-88

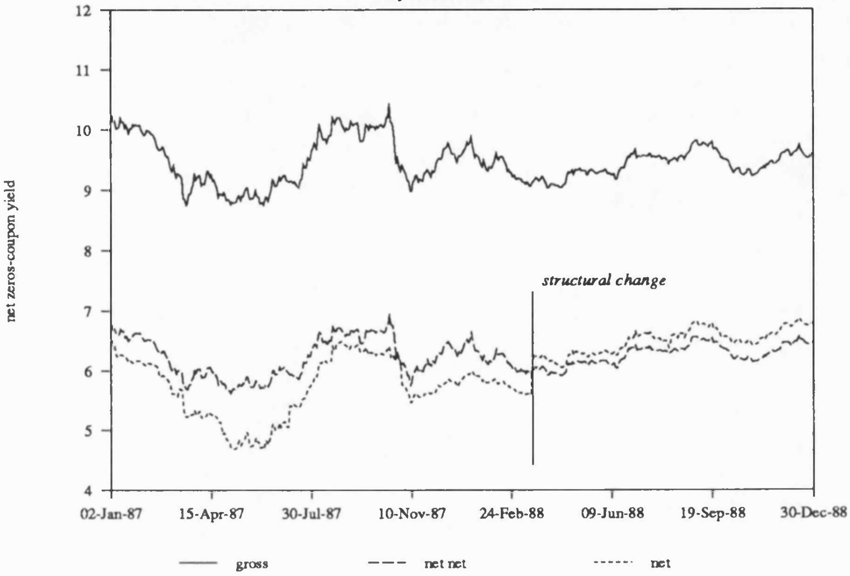

Long (20 year) net yield

daily data 1987-88

The net segment's term structure is, clearly, affected by the change in the rate of income tax used in the model, which was implemented at the time of the 1988 budget. However, even neglecting this, the yield spread between the net and the net net segments' term structures is interesting.

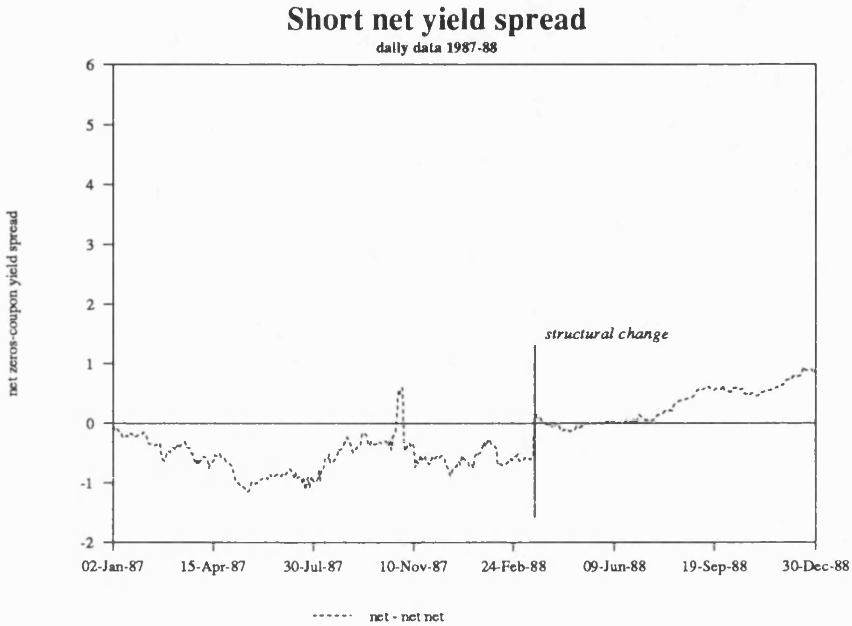

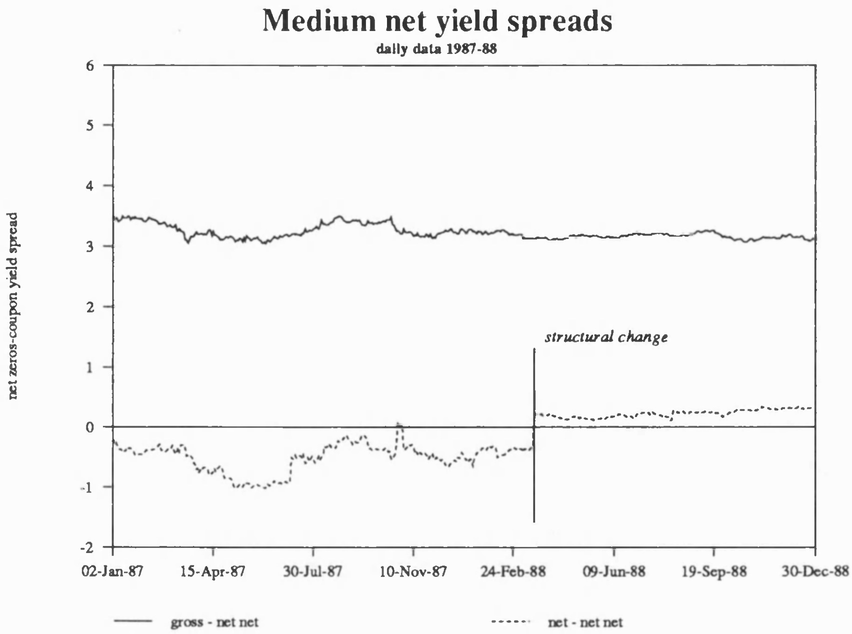

The following table may provide a clue to certain features of this behaviour. It is apparent that a large amount of net segment stock was redeemed in the first half of 1987, and was matched by new issuance with a few months delay, thus probably causing the phase of temporary dearness. It is also interesting to note that the net segment reacted to the stock market crash less violently than the others, as might be expected in the case of a less actively traded segment of the market. Finally, it is worth noting that the relative drift upwards of the net segment's term structure during 1988 is probably due to demand forces, as the personal sector's involvement in the gilt market declined without a corresponding decline in supply (see previous tables).[1]

Table 25. Issues and redemptions of net segment stocks during 1987

date	stock	£ m issued (+) or redeemed (-)
24-Feb-87	Ex 2.5 24-Feb-87	-900
1-May-87	Fn 6.5 1-May-87	-560
13-May-87 †	Tr 8 13-Apr-92	+1,000
14-Jul-87	Tr 3 14-Jul-87	-950
15-Jul-87	Tr 3 11-Jun-92	+500
22-Jul-87 †	Tr 8 10-Dec-91	+1,200

Note: † = stocks were partly paid, and therefore would not have attracted much personal sector interest at first, i.e. until fully paid (approximately 5 weeks later).

1. At the time, personal sector holdings of bank and building society deposits rose sharply, encouraged by high interest rates.

6.3 The covariance of gilt returns

The discount function model adopted has strong implications for the covariance of returns. To be more precise, if the prices are generated by a model of the form:

$$p_t = p(\gamma_t, t) + \varepsilon_t$$

where p_t is the vector of bond prices at time t, γ_t is the vector of model parameters, and ε_t is the vector of errors, the probability distribution of p_t is dependent on that of γ_t and ε_t. Empirical calculations suggest that the error term ε_t contributes only about 1% of the annual stock price volatility, and can therefore probably be neglected for practical purposes.[1] One is therefore in effect left with a n-factor model.

Thus the following approximation should apply to the covariance matrix of gilt returns:

$$\mathrm{cov}(\Delta \ln p_t) \approx J_t \mathrm{cov}(\Delta\gamma_t) J_t'$$

where:

$$J_t = [\partial \ln p(\gamma, t)/\partial\gamma]_{\gamma=\gamma_{t-1}}.$$

Incidentally, this fact might provide a means of estimating the covariance matrix of gilt returns if one were ready to assume that the

1. Furthermore, since errors are broadly independent across stocks, they tend to cancel out in a diversified portfolio.

process γ_t is stationary. Empirical comparisons of this method with other methods commonly used[1] suggests that this might be a sensible methodology, and that this topic probably deserves further research.

The remainder of this section focuses briefly on the main characteristics of the covariance of the term structures, in the hope that these may provide further insight into the behaviour of the factors behind them. The following graph illustrates the volatility of the term structure implied by the model, and shows the well-known tendency of long term rates to be less volatile than short term ones.

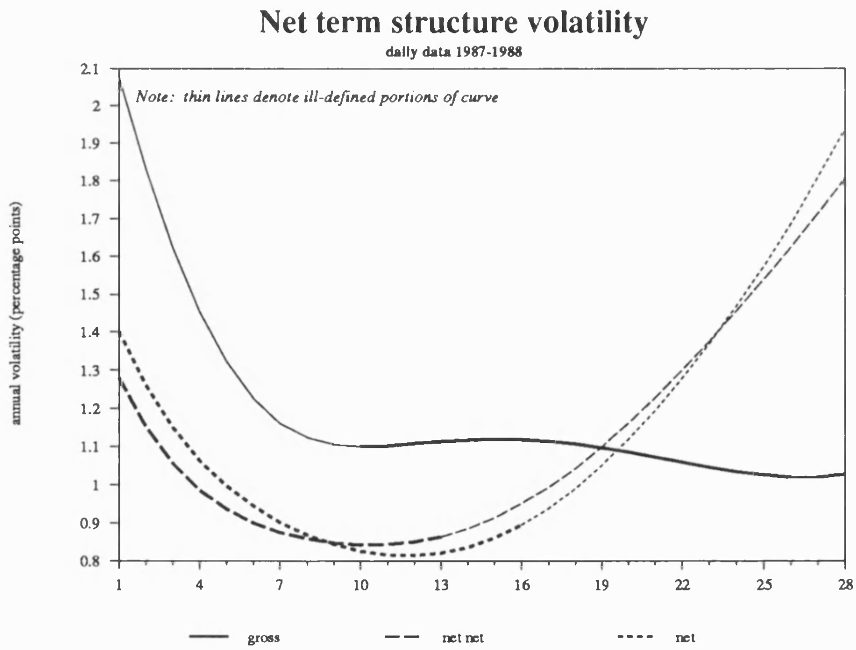

1. Alternative methods are:

i) historical price method: this has the disadvantage that stationarity is usually unacceptable in the case of dated bonds

ii) historical yield method: this has a disadvantage similar to that which afflicts the previous method, although the stationarity assumption is in this case less far-fetched

iii) duration based methods: the approach suggested above could be included in this category, which encompasses a wide range of methodologies.

The following two tables provide indications as to the correlation structure of the process.

Table 26. Correlation matrix daily data 1987-88						
11yr gross	1	0.93	0.92	0.93	0.20	0.59
20yr gross	0.93	1	0.79	0.80	0.15	0.53
4yr net net	0.92	0.79	1	0.92	0.26	0.60
11yr net net	0.93	0.80	0.92	1	0.22	0.57
4yr net	0.20	0.15	0.26	0.22	1	0.65
11yr net	0.59	0.53	0.60	0.57	0.65	1

It is apparent that the correlation between segments is considerable, but that it is usually much less than within segments. This suggests that segment specific factors play a major role in determining movements in the market.

Table 27. Autocorrelation matrix (lag 1) daily data 1987-88						
11yr gross	0.04	0.03	0.07	0.02	0.06	0.08
20yr gross	-0.11	-0.15	-0.06	-0.1	-0.01	-0.03
4yr net net	0.11	0.10	0.10	0.05	0.09	0.12
11yr net net	0.13	0.14	0.11	0.06	0.07	0.12
4yr net	0.21	0.19	0.21	0.20	0.24	0.27
11yr net	0.09	0.09	0.09	0.06	0.11	0.13

There generally appears to positive autocorrelation.[1] Long maturities are an exception, and there is no obvious explanation for this. The

1. At 95% almost all the coefficients in this table are significant (the critical level is about 0.06).

higher order autocorrelations decrease very rapidly and are almost all insignificant as from lag 2. It should be born in mind that the daily price volatility of, say, a medium maturity gilt is of an order of magnitude of about 10 bid-offer spreads, and that if on average 10% of the change on a day carried over to the next, this would still only explain a proportion of volatility of the same order of magnitude as the bid-offer spread. Therefore, while puzzling, a relatively high degree of autocorrelation on a daily basis is not perhaps too unreasonable.

6.4 Conclusions

This chapter has attempted to provoke some thought as to how a model of the demand for gilts might be constructed utilizing the results from a segmented discount function model. Identification of the variables determining the various demand schedules is likely to be complex and require a longer run of data than has been used in this study, and has not therefore been attempted.

The supply component of the market is probably best taken to be exogenous. As described in the next chapter, shocks due to the announcement of individual issues cause short term ripples on the surface, but probably do not significantly affect the term structures. A major advantage of using these term structures would be precisely that they provide a sufficiently detailed description of the price structure to capture micro-economic effects, without suffering from the sort of erratic fluctuations that might afflict individual stocks.

VII. UK GILTS: ANNOUNCEMENT EFFECTS

7.1 Gilt issues

Gilts are issued to the public by minimum price tender, bid price auction and tap. Tenders and auctions are used for the issue of new stock or of large "tranches" of existing stock. The Bank of England issues gilts "on tap" in the case of small "tranchettes" of existing stock that has been created and immediately placed on the Bank's books, or of stock allotted to the Bank at the outcome of a tender or an auction. Since the remainder of this chapter deals with either tender or auction issues, a brief description of the two methods is given below.

Minimum price tenders were the standard method of issue for large amounts of stock before the introduction in 1987 of bid price auctions. Whilst technically quite different, both methods consist of a public invitation to treat. In tenders, bids are called for subject to a minimum tender price.[1] Stock is then allotted at the minimum accepted price. Any stock that is left unsold is picked up by the Bank of England at the minimum accepted price, and later resold on tap. The Bank may also intervene to buy stock at any higher price that it deems appropriate.

Bid price auctions are announced a few weeks before the bid date, and the final details of the issue are made public a few days in advance, whereupon trading starts on a "when issued" basis.[2] There is no minimum tender price. The stock is allotted at the accepted bid price or at the average accepted price depending on the kind of tender

1. Except in the case of index-linked gilts.

2. I.e. for settlement and delivery on the first working day after the auction.

submitted, rather than at the minimum accepted price. As in the case of a minimum price tender, the Bank of England may intervene to buy stock.

7.2 Announcement of issue and exhaustion

Dealers and investors usually have reasonably accurate expectations as to the amount of gilts that will be issued over a period of months. However, they have only a vague idea of the exact characteristics and timing of the issues. Except in the case of the recently introduced auctions, there is little more to go by than the observation that the authorities tend to issue stock in a rising or stable market, and often seek to plug gaps in the spectrum of gilts available.

The announcement of an issue might, therefore, be expected to produce only modest "global effects", i.e. changes in the level of the market as a whole or in the shape of the yield surface. These effects are probably difficult to detect, since they are likely to be small compared to the noise in which they are immersed. By contrast, it is probably more easy to handle the "local effects", i.e. imperfections or ripples in the yield surface affecting stocks of similar characteristics and for which, therefore, the new issue represents a close substitute.

The exhaustion of an issue[1] is also important, since it provides information about the strength of demand, and might be expected to have similar but opposite and presumably smaller effects compared to the announcement of the issue.

1. This will occur either because the issue was entirely bid for at an auction or tender, or because the Bank of England has entirely sold on tap the stock allotted to it.

7.3 Measuring "local" announcement effects

The method used to try and detect ripples on the yield surface is to examine time-series of daily mispricings against the model for a selection of gilts with characteristics similar to those of the newly issued gilt, including the existing stock in the case of a new tranche, spanning the period affected by the issue.

The announcement of an issue might be expected to cause a drop in the price of similar stocks, except in the case of an issue which was widely expected. However, even if an issue were expected, the exact characteristics of the stock can usually only be roughly guessed at except in the case of an auction, when much more information is available in advance, both relating to the characteristics of the stock and the timing of the issue. This observation is supported by the limited empirical evidence available to date. The conclusion might be that auctions are a more efficient way of issuing stock.

One might also suppose that when an issue consists of a tranche of an existing stock, the result may differ from that obtained when the issue consists of a new gilt. The price of the existing stock may sometimes rise because of an improvement in liquidity. The evidence available suggests that this could be the case.

It is reasonable to assume that the larger the issue, the larger the effects. Since the effects are generally small, and only large issues have been considered, it is impossible to confirm this.

Finally, one might surmise that in general the longer the interval between announcement and exhaustion, the smaller the effect of the latter, since it would then be more likely to be expected. The evidence provided suggests that this could be the case.

7.4 Empirical evidence

The period covered by this empirical investigation spans the whole of 1987-88. As previously mentioned, only major issues are considered, i.e. of £1bn or more, with the exception of the second and fourth gilt auctions which were for £900m and £750m respectively. This yields a total of 13 issues, of which 9 were minimum price tenders and 4 were bid price auctions. The more relevant mispricing data is illustrated in the following graphs, where the label *a* denotes issue announcement and *e* denotes exhaustion.

Treasury 10% 9-Jun-94

Most gilts of similar maturity appear to have been affected on the day of the announcement and on the following day. In the cases illustrated below the total fall in mispricing was in the region of 10-30 pence. When the new issue was exhausted, prices tended to bounce back.

14.5 Tr 1-Mar-94

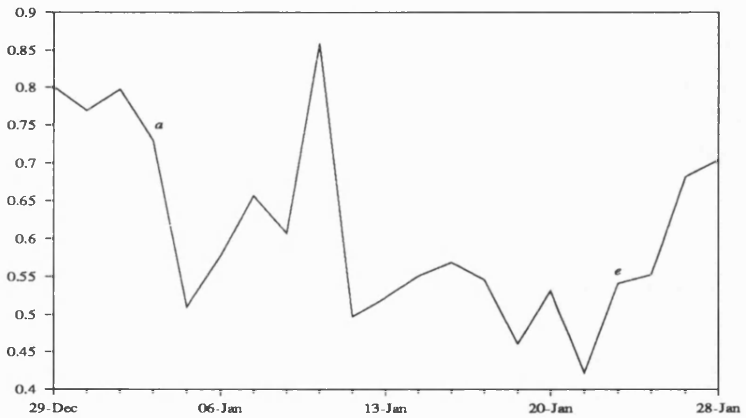

13.5 Ex 27-Apr-94

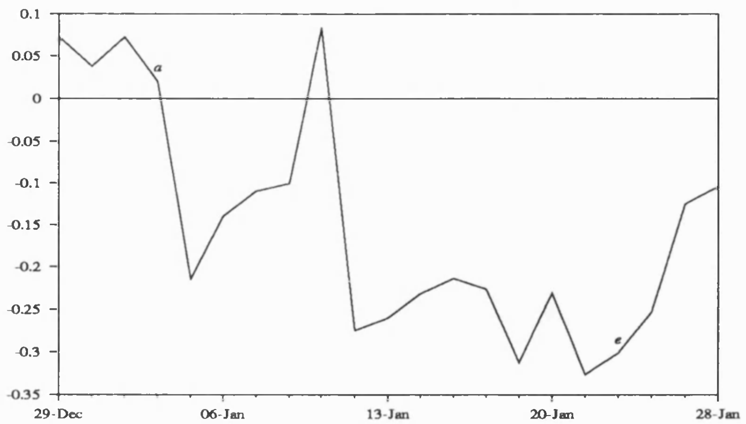

12.5 Ex 22-Aug-94

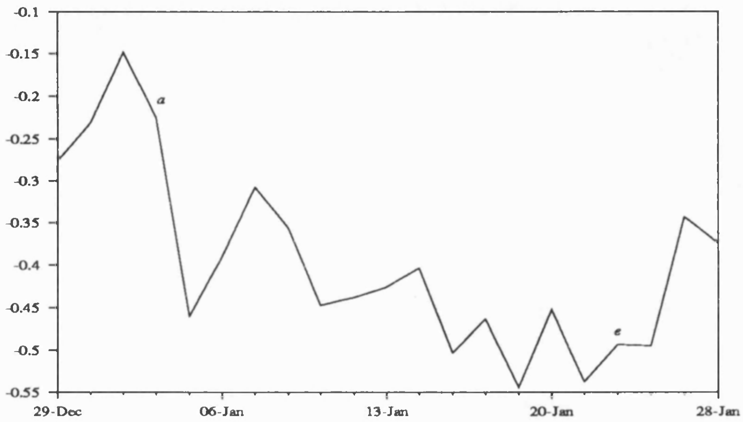

Treasury 9% 13-Oct-08

There was no obvious effect on gilts of similar maturity, but the stock that has the closest duration, i.e. Treasury 8% 2002/06, shows a pattern of behaviour similar to that observed for the previous issue. This could be due to the fact that duration is more important than maturity to investors, and this shows up especially in the case of the longer maturities, for which life and duration differ most.

8 Tr 5-Oct-02/06

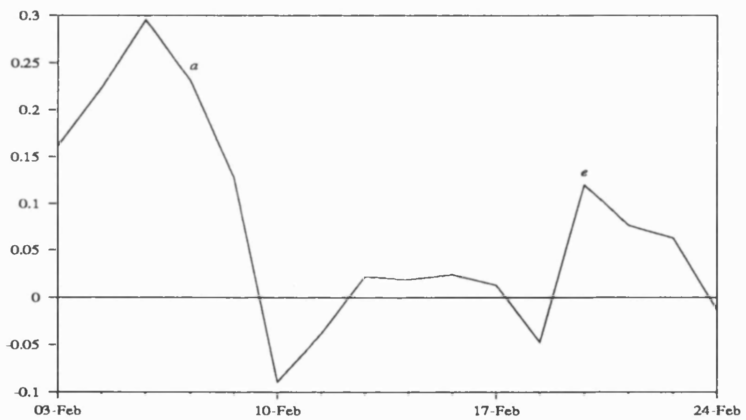

Exchequer 9% 19-Nov-02

Several gilts of similar duration display the familiar pattern of behaviour, but the effect is not very pronounced. The following two examples illustrate the point.

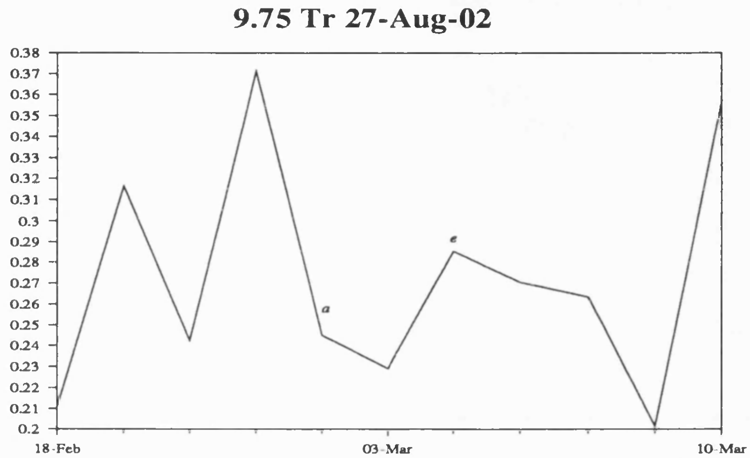

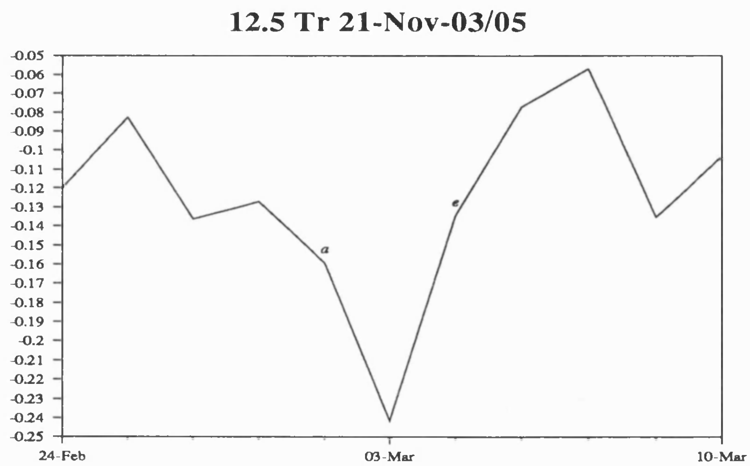

Treasury 8.75% 1-Sep-97 (second tranche)

This issue provides a completely different situation. The day of the announcement was a hectic day, with the announcement both of the tender and of a base rate cut. In the confusion, many stocks moved heavily out of line with the rest of the market. A sharp rise in the mispricing of the existing tranche of the issue on the day of the announcement might also have reflected the expectation of a liquidity effect of the new tranche on the existing stock; the rise was followed by a sharp downward correction, but the overall effect is unclear. Other stocks of similar maturity showed no clear response.

8.75 Tr 1-Sep-97

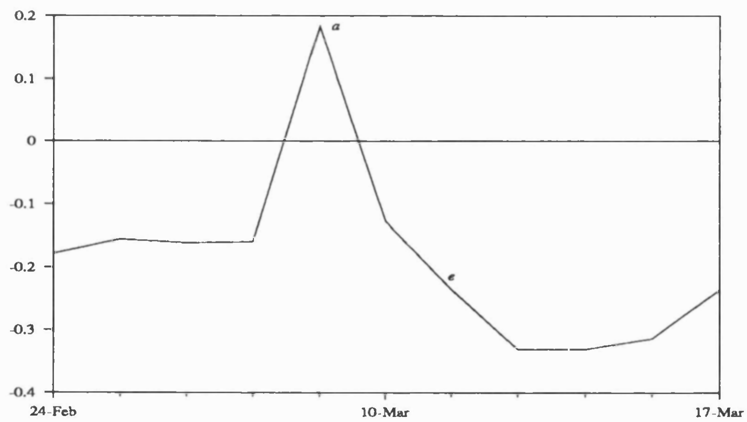

Treasury 8.5% 28-Jan-00

There were few stocks directly competing with this issue. Most of these show a clear response pattern on both announcement and exhaustion dates. The two stocks of closest maturity are provided as examples.

10.25 Cv 22-Nov-99

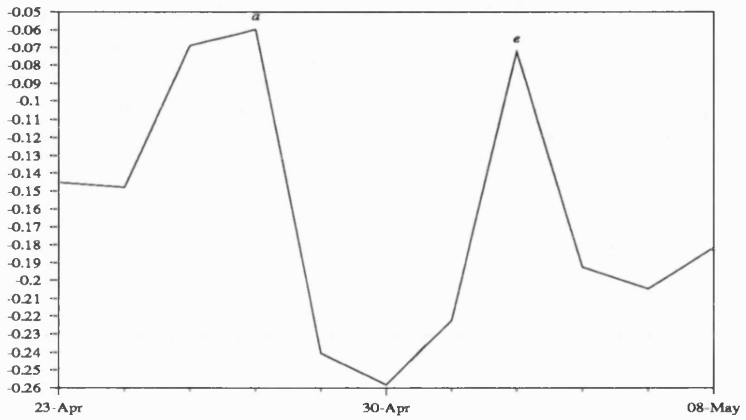

9 Cv 3-Mar-00

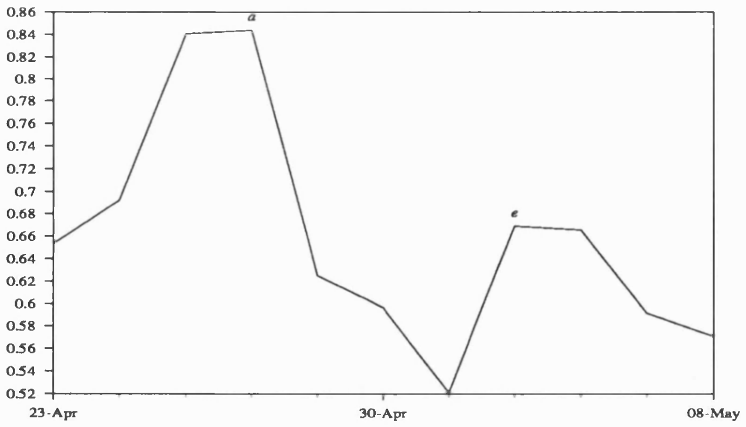

Treasury 8% 13-Apr-92

This was the first auction issue. No response is apparent in gilts of similar characteristics. There were, however, no directly competing stocks. The gap in the spectrum of gilts has since been filled with the issue of £1bn of Treasury 8% 10-Dec-91 and £1bn of Treasury 8% 16-Jul-90. As mentioned previously, the method of issue may also have a bearing on this case.

Treasury 8% 5-Oct-02/06 (second tranche)

There is a clear response pattern on the part of several similar stocks at the announcement of the issue. Much as in the case of Treasury 8.75% 1-Sep-97, the reaction of the original tranche was mooted. Exhaustion, which occurred long after the announcement, appears to have had no effect.

8 Tr 5-Oct-02/06

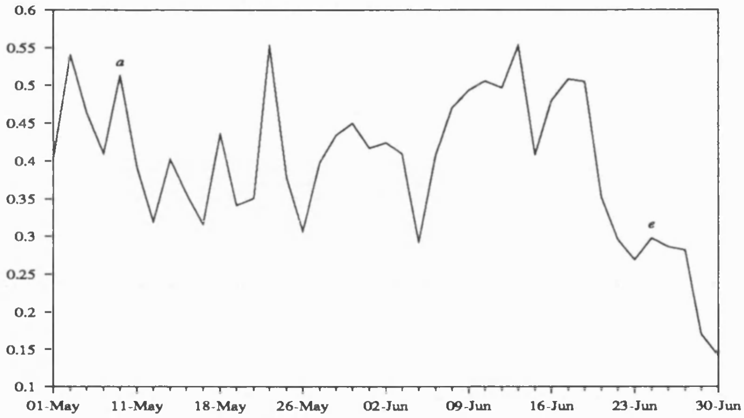

9.75 Cv 15-Nov-06

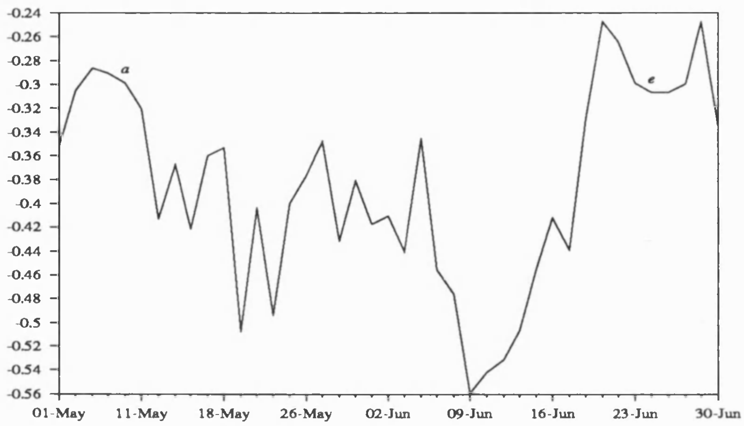

8.5 Tr 16-Jul-07

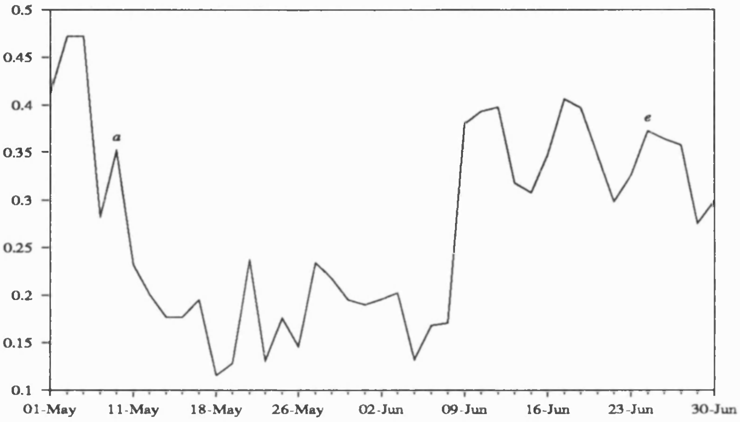

Treasury 8% 10-Dec-91

For several similar stocks there is a clear pattern of response on announcement of the issue. As in the previous case exhaustion, which was much delayed, had no apparent effect.

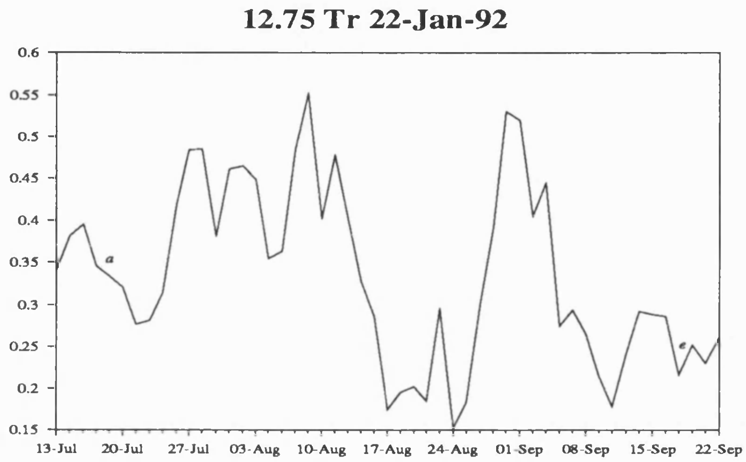

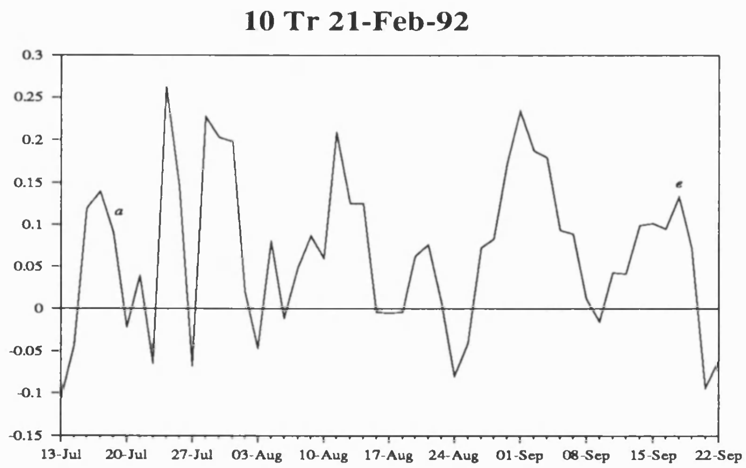

8 Tr 13-Apr-92

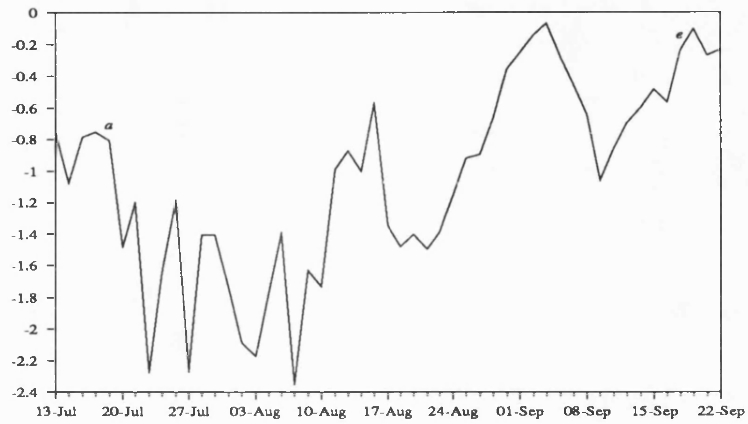

0
-0.2
-0.4
-0.6
-0.8
-1
-1.2
-1.4
-1.6
-1.8
-2
-2.2
-2.4
a
e
13-Jul
20-Jul
27-Jul
03-Aug
10-Aug
17-Aug
24-Aug
01-Sep
08-Sep
15-Sep
22-Sep

Treasury 9% 13-Oct-08 (second tranche)

This was the second auction issue. As in the case of the first auction there is no discernible response. However, this issue was also smaller than the others, amounting to only £900m.

Treasury Convertible 8% 16-Jul-90

The effect of this issue could have been spoilt by the fact that, since it was convertible and the option was at the money, it might have appealed to quite diverse investors. In the circumstance, the issue was treated by the market as a short gilt. As such, the most similar stock is probably Treasury 8% 10-Dec-91. The familiar pattern of response to announcement and exhaustion is again recognizable.

8 Tr 10-Dec-91

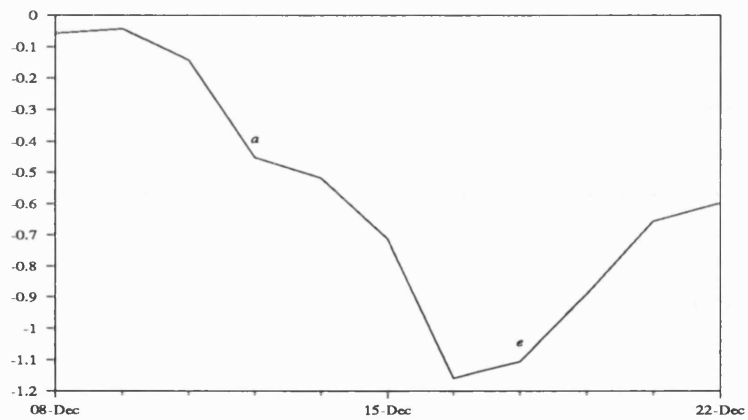

Treasury 8.75% 1-Sep-97 (third tranche)

This was the last of the three experimental auction issues. As in the case of the previous auctions, there appears to have been no systematic response pattern to announcement and exhaustion on the part of similar stocks.

Treasury 8.5% 3-Feb-94

There are no stocks directly competing with this issue. No clear pattern of response to announcement and exhaustion was detected.

Treasury 8.5% 3-Feb-94 (second tranche)

This was the fourth auction issue. An albeit somewhat vague response pattern is apparent. This may be related to the fact that, differently from the previous auctions, the choice of gilt took the market completely by surprise.

8.5 Tr 3-Feb-94

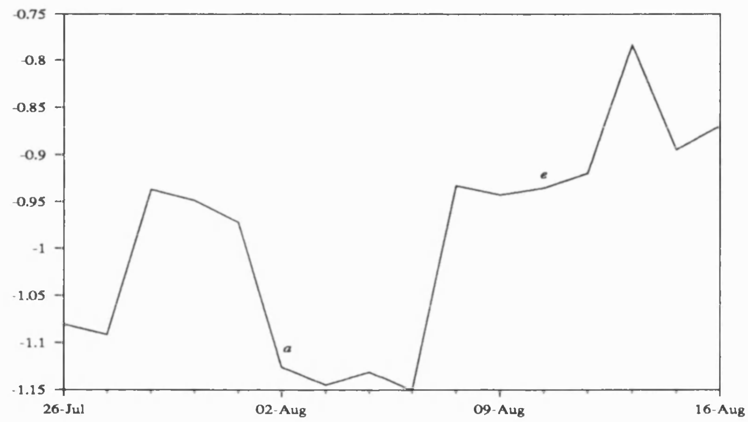

14.5 Tr 1-Mar-94

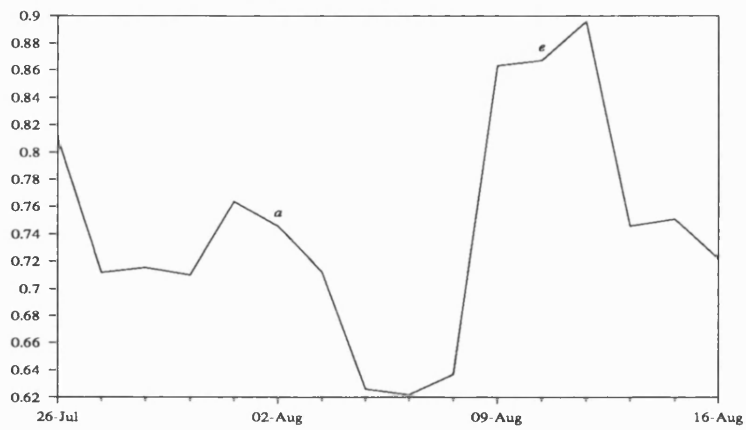

7.5 Final remarks

The effects of announcement and exhaustion are fairly clear in the case of most tender issues, and follow reasonable patterns. However, the size of the effects is generally small, more often than not around 10-30 pence for the stock most affected, although occasionally as large as 50 pence. The effects are therefore usually only about 2 or 3 times the bid-offer spread, suggesting that the gilt market is rather efficient. It also appears that effects tend to be spread over several days, which is a curious, albeit very small, type of imperfection.

These points are summarised by the following graph, which shows the average pattern of response for the most affected stock in each tender.

Average error effect pattern

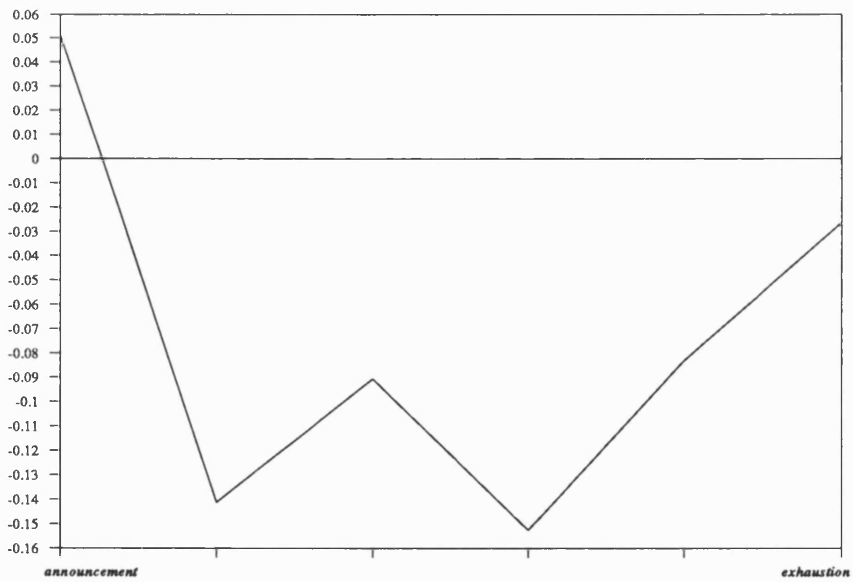

CONCLUSIONS

This study advocates the use of a new specification of discount function model which is easy to compute, well-behaved, and widely applicable.

After defining formal conditions under which discount function models are suitable, the study goes on to apply a basic version of the proposed specification in the case of two government bond markets that are reputedly free from segmentation and taxation effects. The results are broadly satisfactory. In particular, the standard error of either model is small, i.e. of the same order of magnitude as the bid-offer spread. The main weakness of this method appears to be the inability to capture liquidity effects.

The basic discount function model is then extended to cope with segmentation and taxation. The solution proposed is to estimate simultaneously discount functions for each of a set of representative investors, adjusting for tax payments in the case of taxation. After highlighting the weaknesses inherent in net of tax models, the study proceeds to apply the proposed model to the UK gilt market. The results are not as satisfactory as in the cases previously examined. In particular, the standard error is over 3 times the bid-offer spread. A major factor contributing to this is identified in the FOTRA effect.

The implicit tax rates obtained from the UK gilt model estimation appear to agree with those in force for the major investor types. In particular, it is shown how, in the run up to the income tax changes implemented with the 1988 budget, the implicit tax rates changed in response to expectations of a tax cut.

Given the diversity in liquidity that exists in the gilt market between stocks of similar coupon and maturity, it was decided to try and identify factors affecting individual stock standard errors. The latter are shown to be related to the bid-offer spread rather than duration, suggesting that one of the main reasons adduced for modelling bond markets in a yield space, i.e. homoscedasticity in yield terms, is invalid. However, it is shown that allowing for heteroscedasticity does not in the case of the gilt market produce significantly different results.

The tax-timing option described by Constantinides and others has some fairly robust implications for the behaviour of the model residuals of gilts with high coupons and short to medium maturities. Empirical evidence, however, suggests that the gilt market place does not recognize anything like the option value that the theory would seem to predict.

The estimated gilt term structures are shown to react in a sensible way to changes in demand and supply. The effect of the announcement of a new issue on adjacent stocks is also clearly identified by the model, although it appears to be on average only about 2 or 3 times the size of the bid-offer spread.

On the basis of the evidence collected in this study, it would appear that more work on simultaneous estimation of segmented discount functions would be fruitful. In particular, this might be worth attempting in the case of segmentation without taxation. It would also probably be worth experimenting with a larger number of segments in the case of the gilt market.

Finally, identifying and estimating individual segment demand schedules, while complicated, could be an extremely valuable endeavour.

APPENDIX[1]

A.1 Coupon effects

A.1.1 Duration effect

The fundamental identity defining a yield in case of no taxation or segmentation is:

$$q(y,b) = p(v,b).$$

This can be re-written as:

$$b'w(y) = b'v$$

where $w(y)$ is the discount function implied by the yield y, i.e. the vector with n-th component equal to $(1 + y/100)^{-n}$.

Since $b = c\,u + 100\,e$, where u denotes an appropriate vector of ones and noughts, and e denotes the appropriate unit vector, one gets:

$$c\,u'[v - w(y)] + 100\,e'[v - w(y)] = 0.$$

1. The notation used throughout this section is explained in chapter I.

Taking $c = 0$ and $c \to \infty$ in this expression it is easily seen that:

$$\begin{cases} e'[v - w(y_0)] = 0 \\ u'[v - w(y_\infty)] = 0 \end{cases} \Rightarrow \begin{cases} e'v = e'w(y_0) \\ u'v = u'w(y_\infty) \end{cases}$$

where y_0 and y_∞ denote the zero-coupon and the annuity yield.

By substitution:

$$c\, u'[w(y_\infty) - w(y)] + 100\, e'[w(y_0) - w(y)] = 0$$

and therefore, since each component of $w(y)$ is monotonically decreasing in y:

$$\begin{cases} y > y_0 \Leftrightarrow y_\infty > y \Leftrightarrow y_\infty > y_0 \\ y < y_0 \Leftrightarrow y_\infty < y \Leftrightarrow y_\infty < y_0. \end{cases}$$

Totally differentiating gives:

$$u'[w(y_\infty) - w(y)]dc + b'[\partial w(y)/\partial y]dy = 0$$

and since $b'[\partial w(y)/\partial y] < 0$, one concludes that:

$$\begin{cases} dy/dc > 0 \Leftrightarrow y_\infty > y_0 \\ dy/dc < 0 \Leftrightarrow y_\infty < y_0. \end{cases}$$

The coupon effect can be related to the annuity yield curve in a simple way by considering that, in obvious notation:

$$\sum_{\tau=1}^{t} w_\tau(y_{\infty,t+1}) + w_{t+1}(y_{\infty,t+1}) =$$

$$\sum_{\tau=1}^{t} v_\tau + v_{t+1} =$$

$$\sum_{\tau=1}^{t} w_\tau(y_{\infty,t}) + w_{t+1}(y_{0,t+1})$$

so:

$$\sum_{\tau=1}^{t} [w_\tau(y_{\infty,t+1}) - w_\tau(y_{\infty,t})] + [w_{t+1}(y_{\infty,t+1}) - w_{t+1}(y_{0,t+1})] = 0.$$

Therefore:

$$\begin{cases} y_{\infty,t+1} > y_{\infty,t} \Leftrightarrow y_{0,t+1} > y_{\infty,t+1} \\ y_{\infty,t+1} < y_{\infty,t} \Leftrightarrow y_{0,t+1} < y_{\infty,t+1} \end{cases}$$

i.e the coupon effect is of opposite sign to that of the slope of the annuity yield curve.

A.1.2 Tax effect

In the case of a taxed investor in gilts who buys and holds to maturity, the fundamental identity becomes:

$$q(y,b) = p(v,b,\alpha,\gamma).$$

This can be re-written as:

$$b'\,w(y) = \bar{b}'v$$

or equivalently:

$$c\,[\bar{u}'v - u'\,w(y)] + 100\,[\bar{e}'v - e'w(y)] = 0$$

where:

$$\bar{b} = c\,\bar{u} + 100\,\bar{e}$$

$$\bar{u} = [(1-\alpha)/(1-\gamma v_t)]\,u$$

$$\bar{e} = [(1-\gamma)/(1-\gamma v_t)]\,e.$$

Taking $c = 0$ and $c \to \infty$ it is seen that:

$$\begin{cases} \bar{e}'v = e'w(y_0) \\ \bar{u}'v = u'w(y_\infty) \end{cases}$$

in analogy with the no-tax case.

By substitution:

$$c\,u'[w(y_\infty) - w(y)] + 100\,e'[w(y_0) - w(y)] = 0$$

which is formally identical to the expression obtained in the no-tax case. Thus one concludes that:

$$\begin{cases} dy/dc > 0 \Leftrightarrow y > y_0 \Leftrightarrow y_\infty > y \Leftrightarrow y_\infty > y_0 \\ dy/dc < 0 \Leftrightarrow y < y_0 \Leftrightarrow y_\infty < y \Leftrightarrow y_\infty < y_0. \end{cases}$$

The coupon effect depends on the tax rates and on the discount function. For example, given the gross discount factor curve z, i.e the vector with n-th component equal to $[(1-\gamma)/(1-\gamma v_t)]v_t$, differentiating the price function:

$$p = p(v(z,\gamma), b, \alpha, \gamma)$$

would give:

$$\partial p/\partial\alpha = [\partial\bar{b}/\partial\alpha]'v = -\,c\,u'v/(1-\gamma v_t) < 0$$

i.e. the gross yield would rise with the rate of income tax.

Similarly, if the par yield curve k were given, differentiating the price function:

$$p = p(v(k,\alpha), b, \alpha, \gamma)$$

would give:

$$\partial p/\partial\gamma = [\partial\bar{b}/\partial\gamma]'v = [(p-100)v_t]/(1-\gamma v_t)$$

i.e. above par the gross yield falls with the rate of tax on capital gains, while below par it rises. So, as one would expect, the taxation of capital gains and of income have, in a sense, opposite effects on the shape of the yield surface.

A.2 Immunization and discount functions

Immunization[1] and discount function theory are in general incompatible.[2] In fact, within a discount function model, a parallel shift in the yield surface is possible only if the surface is flat. Furthermore, only if the surface is flat can all points on the surface of equal duration have equal yields. These two facts result from the following theorems. Whilst fairly intuitive, neither of these results is, to my knowledge, available in the literature.

---

1. Immunization is here used in the sense of the traditional duration-based bond portfolio management technique. This is, essentially, a method of optimization which assumes parallel shifts in the instantaneous yield surface. An important consequence of this model of investor behaviour is that any two bonds with the same duration should have the same instantaneous yield. Similar results are obtained if modified duration and n-annual yield are substituted for duration and instantaneous yield.

2. With no meaningful loss of generality, only the case of no taxation or segmentation is considered here.

Theorem 1

The yield surface can shift in parallel if and only if it is flat.

Proof

The fact that a parallel shift is possible if the surface is flat is obvious. The converse is proved hereafter.

The fundamental identity defining a yield is:

$q(\bar{y}, b) = p(v,b).$

Substituting appropriately and totally differentiating, bearing in mind that $[\partial\, w(\bar{y})/\partial \bar{y}] = \Lambda\, w(\bar{y})$ where Λ denotes the diagonal matrix with term $\lambda_{\tau,\tau} = \tau$, and that in the case of a parallel shift $d\bar{y} = d\bar{y}_\infty = d\bar{y}_0$, one gets:

$$\begin{cases} c\, u'\, w(\bar{y}) + 100\, e'\, w(\bar{y}) = c\, u'\, w(\bar{y}_\infty) + 100\, e'\, w(\bar{y}_0) \\ c\, u'\Lambda\, w(\bar{y}) + 100\, e'\Lambda\, w(\bar{y}) = c\, u'\Lambda\, w(\bar{y}_\infty) + 100\, e'\Lambda\, w(\bar{y}_0). \end{cases}$$

Subtracting the second equation divided by t from the first gives:

$$c\, u'(I - \ddot{\Lambda})\, [w(\bar{y}) - w(\bar{y}_\infty)] = 0$$

where $\ddot{\Lambda} = \Lambda/t.$

Clearly the equation is satisfied if $c = 0$, or $\bar{y} = \bar{y}_\infty$, or $t = 1$. No other solution is possible since this would require:

$$u'(I - \ddot{\Lambda})\, [w(\bar{y}) - w(\bar{y}_\infty)] = 0$$

with $u'(I - \ddot{\Lambda}) > 0$ and either $w(\bar{y}) > w(\bar{y}_\infty)$ or $w(\bar{y}) < w(\bar{y}_\infty)$, which is impossible.

For $t = 1$ there never is a coupon effect. For $t > 1$ there cannot be either, since then $c > 0 \Rightarrow \bar{y} = \bar{y}_{\infty}$, and an immediate result of the analysis of section A.1.1 is that then $\bar{y} = \bar{y}_{\infty} = \bar{y}_0$ all c.

But, by a further application of the results of section A.1.1, if there is no coupon effect at any maturity, the annuity yield curve, and hence the whole yield surface, must be flat. ■

Theorem 2

All points on the yield surface of equal duration have equal yield if and only if the surface is flat.

Proof (outline)

The fact that all point on the surface of equal duration have equal yield if the surface is flat is obvious. The converse is not.

From the discussion of section A.1.1, the yield is a monotonic function of coupon, and takes, by continuity, all values from $\bar{y}_0$ to $\bar{y}_\infty$.

Since duration is given by:

$$d = b'\Lambda\, w(\bar{y})/b'\, w(\bar{y})$$

one finds, chosing $c = 0$ and taking the limit for $c \to \infty$, that:

$$\begin{cases} d_0 = e'\Lambda\, w(\bar{y}_0)/e'\, w(\bar{y}_0) = t \\ d_\infty = u'\Lambda\, w(\bar{y}_\infty)/u'\, w(\bar{y}_\infty). \end{cases}$$

For $t > 2$ it is easily shown that $d_\infty < (t+1)/2 \leq t-1$. Therefore, by continuity, duration takes all values in the non-trivial interval $(d_\infty, t-1)$. It follows that one can choose a pair of bonds for each of the maturities $t, t-1 \geq 2$ so that they share the same pair of durations in the interval $(d_\infty, t-1)$. If there were no coupon effect at one maturity, the relevant pair would have the same yield, and if all bonds with the same duration were assumed in general to have the same yield, the other pair would also have to have same yield, thus ensuring that there could be no coupon effect at that maturity either. By induction, one concludes that, if there were no coupon effect at a maturity $t > 1$, the whole yield surface would have to be flat.[1]

Suppose that the surface were not flat, so that for any maturity $t > 1$ there were a coupon effect. The fundamental identity defining a yield could then be re-written as:

$$c = 100\, e'[w(\bar{y}) - v]/u'[w(\bar{y}) - v]$$

and substituting this into the expression for duration one would get:

$$d = \frac{u'[w(\bar{y}) - v]e'\Lambda w(\bar{y}) + e'[w(\bar{y}) - v]u'\Lambda w(\bar{y})}{u'[w(\bar{y}) - v]e'w(\bar{y}) + e'[w(\bar{y}) - v]u'w(\bar{y})}.$$

Applying this expression to two maturities $t, t-1 \geq 2$, on the assumption that bonds of equal duration must have the same yield, then:

$$\frac{u'_1[w(\bar{y}) - v]e'_1\Lambda w(\bar{y}) + e'_1[w(\bar{y}) - v]u'_1\Lambda w(\bar{y})}{u'_1[w(\bar{y}) - v]e'_1 w(\bar{y}) + e'_1[w(\bar{y}) - v]u'_1 w(\bar{y})} =$$

$$= \frac{u'_2[w(\bar{y}) - v]e'_2\Lambda w(\bar{y}) + e'_2[w(\bar{y}) - v]u'_2\Lambda w(\bar{y})}{u'_2[w(\bar{y}) - v]e'_2 w(\bar{y}) + e'_2[w(\bar{y}) - v]u'_2 w(\bar{y})}$$

for any yield $\bar{y}$ in the non-trivial interval corresponding to $(d_\infty, t-1)$.

But the two sides of the equation are not identically equal, so the assumption that the yield surface wasn't flat cannot be right. ■

1. The surface is assumed to be defined at annually spaced maturities. In the case of n-annual bonds it may be desirable to define the surface at n-annually spaced maturities. The proof can then be modified accordingly.

BIBLIOGRAPHY

1. **Bank of England** (Mathematical Techniques Group, Financial Statistics Division) "Yield curves for gilt-edged stocks: an improved model" *Bank of England Quarterly Bulletin*, June 1982.

2. **Bank of England** (Mathematical Techniques Group, Financial Statistics Division) "A new yield curve model" *Bank of England Quarterly Bulletin*, February 1990.

3. **Brennan, M. J.** and **E. S. Schwartz** "A continuous time approach to the pricing of bonds" *Journal of Banking and Finance*, July 1979.

4. **Burman, J. P.** "Yield curves for gilt-edged stocks: further investigation" *Bank of England Quarterly Bulletin*, September 1973.

5. **Burman, J. P.** and **O. Page** "Yield curves for gilt-edged stocks: a further modification" *Bank of England Quarterly Bulletin*, June 1976.

6. **Burman, J. P.** and **W. R. White** "Yield curves for gilt-edged stocks" *Bank of England Quarterly Bulletin*, December 1972.

7. **Buse, A.** "Expectations, prices, coupons and yields" *Journal of Finance*, September 1970.

8. **Caks, J.** "The coupon effect on yield to maturity" *Journal of Finance*, March 1977.

9. **Campbell, J. Y.** "A defence of traditional hypotheses about the term structure of interest rates" *Journal of Finance*, March 1986.

10. **Carleton, W. T.** and **I. A. Cooper** "Estimation and uses of the term structure of interest rates" *Journal of Finance*, September 1976.

11. **Carr, J. L., P. J. Halpern** and **J. S. McCallum** "Correcting the yield curve: a re-interpretation of the duration problem" *Journal of Finance*, September 1974.

12. **Chambers, D. R., W. T. Carleton** and **D. W. Waldman** "A new approach to estimation of the term structure of interest rates" *Journal of Financial and Quantitative Analysis*, September 1984.

13. **Clarkson, R. S.** "A mathematical model for the gilt-edged market" *Transactions of the Faculty of Actuaries*, March 1978.

14. **Constantinides, G. M.** "Capital market equilibrium with personal tax" *Econometrica*, May 1983.

15. **Constantinides, G. M.** "Optimal stock trading with personal taxes: implications for prices and the abnormal January returns" *Journal of Finance*, August 1982.

16. **Constantinides, G. M.** and **J. E. Ingersoll jr** "Optimal bond trading with personal taxes" *Journal of Financial Economics*, April 1984.

17. **Constantinides, G. M.** and **M. S. Scholes** "Optimal liquidation of assets in the presence of personal taxes: implications for asset pricing" *Journal of Finance*, May 1980.

18. **Cox, J. C., J. E. Ingersoll jr** and **S. A. Ross** "A re-examination of traditional hypotheses about the term structure of interest rates" *Journal of Finance*, September 1981.

19. **Cox, J. C., J. E. Ingersoll jr** and **S. A. Ross** "A theory of the term structure of interest rates" *Econometrica*, March 1985.

20. **Cox, J. C., J. E. Ingersoll jr** and **S. A. Ross** "An intertemporal general equilibrium model of asset prices" *Econometrica*, March 1985.

21. **Cramer, R.** and **S. L. Hawk** "The consideration of coupon levels, taxes, reinvestment rates, and maturity in the investment management of financial institutions" *Journal of Financial and Quantitative Analysis*, March 1975.

22. **Dammon, R. M., K. B. Dunn** and **C. S. Spatt** "A reexamination of the value of tax options" *Review of Financial Studies*, 1989.

23. **Dermody, J. C.** and **E. Z. Prisman** "Term structure multiplicity and clientele in markets with transaction costs and taxes" *Journal of Finance*, September 1988.

24. **Dobson, S. W.** "Estimating term structure equations with individual bond data" *Journal of Finance*, March 1978.

25. **Echols, M. E.** and **J. W. Elliott** "A quantitative yield curve model for estimating the term structure of interest rates" *Journal of Financial and Quantitative Analysis*, March 1976.

26. **Elliott, J. W.** and **M. E. Echols** "Market segmentation, speculative behaviour, and the term structure of interest rates" *Review of Economics and Statistics*, February 1976.

27. **Fogler, H. R.** and **S. Ganapathy** "Comment on: A quantitative yield curve model for estimating the term of structure of interest rates" *Journal of Financial and Quantitative Analysis*, June 1980.

28. **Hodges, S. D.** and **S. M. Schaefer** "A model for bond portfolio improvement" *Journal of Financial and Quantitative Analysis*, June 1977.

29. **Jordan, J. V.** "Tax effects in term structure estimation" *Journal of Finance*, June 1984.

30. **Khang, C.** "Expectations, prices, coupons and yields: comment" *Journal of Finance*, September 1975.

31. **Litzenberger, R. H.** and **J. Rolfo** "An international study of tax effects on government bonds" *Journal of Finance*, March 1984.

32. **Litzenberger, R. H.** and **J. Rolfo** "Arbitrage pricing, transaction costs and taxation of capital gains: a study of government bonds with the same maturity date" *Journal of Financial Economics*, April 1984.

33. **Livingston, M.** "A theory of humpbacked bond yield curves" *Journal of Finance*, December 1977.

34. **Livingston, M.** "Bond taxation and the shape of the yield to maturity curve" *Journal of Finance*, March 1979.

35. **Livingston, M.** "Flattening of bond yield curves for long maturities" *Journal of Finance*, March 1982.

36. **Livingston, M.** "Taxation and bond market equilibrium in a world of uncertain future interest rates" *Journal of Financial and Quantitative Analysis*, March 1979.

37. **Livingston, M.** "The pricing of premium bonds" *Journal of Financial and Quantitative Analysis*, September 1979.

38. **Livingston, M.** and **J. Caks** "A 'duration' fallacy" *Journal of Finance*, March 1977.

39. **McCallum, J. S.** "The impact of the capital gains tax on bond yields" *National Tax Journal*, December 1973.

40. **McCulloch, J. H.** "Measuring the term structure of interest rates" *Journal of Business*, January 1971.

41. **McCulloch, J. H.** "The tax-adjusted yield curve" *Journal of Finance*, June 1975.

42. **Nelson, C. R.** and **A. F. Siegel** "Parsimonious modelling of yield curves" *Journal of Business*, October 1987.

43. **Robichek, A. A.** and **W. D. Niebuhr** "Tax-induced bias in reported treasury yields" *Journal of Finance*, December 1970.

44. **Ronn, E. I.** "A new linear programming approach to bond portfolio management" *Journal of Financial and Quantitative Analysis*, December 1987.

45. **Schaefer, S. M.** "Measuring a tax-specific term structure of interest rates in the market for British government securities" *Economic Journal*, June 1981.

46. **Schaefer, S. M.** "Tax-induced clientele effects in the market for British government securities" *Journal of Financial Economics*, July 1982.

47. **Schaefer, S. M.** "Taxes and security market equilibrium" *Financial Economics: Essays in Honour of Paul Cootner*, 1982.

48. **Schaefer, S. M.** "The problem with redemption yields" *Financial Analysts Journal*, July 1977.

49. **Shea, G. S.** "Interest rate term structure estimation with exponential splines: a note" *Journal of Finance*, March 1985.

50. **Shea, G. S.** "Pitfalls in smoothing interest rate term structure data: equilibrium models and spline approximations" *Journal of Financial and Quantitative Analysis*, September 1984.

51. **Shiller, R. J.** and **F. Modigliani** "Coupon and tax effects on new and seasoned bond yields and the measurement of the cost of debt capital" *Journal of Financial Economics*, August 1979.

52. **Vasicek, O. A.** and **H. G. Fong** "Term structure modelling using exponential splines" *Journal of Finance*, May 1982.