A Simple Model of the Taxable and Tax-Exempt Yield Curves

Richard C. Green
Carnegie Mellon University

I examine the anomalous behavior of the taxable and tax-exempt yield curves. Long municipal yields appear too high relative to the equivalent after-tax yield that can be earned in Treasury or corporate bonds. I discuss existing explanations of the problem and propose a simple model that relates the yields of taxable bonds to the yield curve for par tax exempts. The ratio of the tax-exempt yield to the taxable yield increases with maturity in the model, so it is consistent with observed phenomena such as inverted yield curves for taxables and contemporaneous rising yield curves for tax exempts. Statistical and descriptive comparisons between the yields predicted by the model and observed yields on par bonds show that the model has some promise in explaining the apparent anomalies in the behaviors of the two yield curves.

The market for municipal bonds provides an excellent laboratory in which to evaluate the impact of

My interest in these issues began as part of a project course sponsored by Miller, Anderson, and Sherrerd. I would like to thank Ken Dunn and Steve Kreider of MA&S for their help and especially for many thoughtful conversations about this analysis. I would like to thank the students in the course for help with preliminary computations and for educating me about the muni market. Part of this work was completed while I was visiting the Stockholm School of Economics. Seminar participants at UBC, UC-Irvine, Caltech, Cornell, HEC, INSEAD, LSE, MIT, Michigan, Wisconsin-Madison, and the 1992 AFA meetings provided many useful comments. In addition, I have benefited greatly from the comments of Bob Dammon, David Goldreich, Burton Hollifield, Bernt Oedegaard, Ehud Ronn, Chester Spatt, Rob Stambaugh (the editor), Suresh Sundaresan, Arthur Warga, and an anonymous referee. Address correspondence to Richard C. Green, GSIA, Carnegie Mellon University, Pittsburgh, PA 15213.

The Review of Financial Studies 1993 Volume 6, number 2, pp. 233-264

taxation on relative asset values. Investors are concerned with the returns they earn after tax. Holding other things constant, one expects the forces of supply and demand to equate after-tax returns. Coupon income on municipals is tax-exempt. On other bonds it is taxed as ordinary income. Thus, one might expect the relative yields of par municipals and par taxable bonds to reflect the tax rate of some investor who could conceivably be indifferent to the two types of bonds.

For short-maturity instruments the data conform reasonably well to these expectations. The standard rule of thumb for evaluating taxable versus tax-exempt bonds compares the tax-exempt yields to the taxable yield times 1 minus the tax rate. Thus, in April of 1992, *Barron's* reported an average seven-day yield of 3.80 percent for the taxable money funds it reports on, while the same measure for the tax-free money market funds was 2.77 percent. At these yields the implied tax rate is 27.11 percent, reasonably close to the marginal rates for high-income individuals and corporations, which are close to one-third. The same issue reports average yields of 6.21 percent and 4.47 percent for one year previously, for an implied tax rate of 28.02 percent *(Barron's*, April 6, 1992, p. 127).

A number of researchers, however, have documented anomalous behavior for long-maturity bonds. The long municipals have yields that seem to be too high, relative to any reasonable marginal tax rate, and this behavior is very persistent across time. The issue of *Barron's* cited gives the average yield on the Bond Buyer Municipal Bond Index as 6.89 percent. Barron's Best Grade Bond index had an average yield of 8.31 percent. This implies a critical tax rate of 17.09 percent, which seems particularly low in light of expectations at this time that tax rates would be rising in the future. The same statistics from the previous year give a critical tax rate of 19.31 percent *(Barron's,* April 6, 1992, p. 147).

In fact, the whole term structure behaves differently for municipals. Fabozzi and Feldstein (1983, pp. 401-403) state:

> *Two characteristics of the municipal yield curve should be noted. First, the municipal yield curve almost always has a positive slope. . . . Second, the yield spreads between maturities are usually wider in the municipal market compared with maturity spreads in the Treasury market.*

Salomon Brothers' *Analytical Record of Yields and Yield Spreads* (1991) provides monthly estimates of current-issue prime-grade municipal yields for various maturities from 1954 through 1990. Over this entire period there has never been a month when the 1-year yield exceeded the 30-year yield. The yield curve for Treasury and corporate securities, on the other hand, is frequently "inverted." During the most extreme of these inversions, in March and April of 1980, one-year Treasuries had yields over 300 basis points above those of 30-year Treasuries. Yield curves for mortgage-backed bonds and corporate bonds were also inverted at this time.[1] Yet 30-year prime-grade municipals yielded 150 and 80 basis points over the one-year tax exempts in those months, respectively. A more thorough discussion of this evidence, along with a review of previous work on this problem, is provided in Section 1.

I develop a simple model that reproduces this behavior. The differences between the approach taken here and the traditional one are fairly intuitive. Traditional analyses of the problem assume that taxable investors are "at the margin" between coupon-bearing taxables and tax exempts of the same maturity. These investors apply the same discount factors (or marginal rates of substitution) to the after-tax cash flows from both types of bonds. If we consider bonds of each type that are priced at par and are default-free, and if the investor plans to hold them to maturity, it follows that the yield on the tax exempt is the yield on the taxable times 1 minus the tax rate. The taxable par yield curve is a constant multiple of the tax-exempt par yield curve, and this appears to be inconsistent with the data.

But why would we expect high-tax investors, who are presumably holding the municipals, to hold par taxables in isolation? If there are portfolios of taxables that are tax-advantaged relative to par bonds but that have similar pretax cash flows, we might more reasonably expect high-tax investors to hold those. In particular, we would anticipate attempts by taxable investors to shield coupon income from taxation by constructing portfolios that generate offsetting losses or investment interest expense. (We assume that investors can do this in markets that, aside from taxation, are quite frictionless. Unlimited tax arbitrage is ruled out by the restrictions in the tax code.) If taxable investors are "marginal" across these portfolios and municipals, they apply the same discount factors to the after-tax cash flows from both. Using this relationship, we can obtain their implicit valuations of the pretax cash flows from par taxable bonds.

How, then, should par taxables be priced relative to these portfolios of taxables that are tax-advantaged? Here we can appeal to the arbitrage activities of dealers and tax-exempt institutions. While tax rules limit their ability to arbitrage across the taxable and tax-exempt yield curves, they are relatively free to arbitrage out any tax-induced pricing differentials within the taxable yield curve itself. Arbitrage by dealers should drive together the prices of par taxables and other portfolios

[1] See Salomon Brothers *Analytical Record of Yields and Yield Spread* (1990), Part II, Table 12A, page 4, and Table 13, page 11.

of taxable bonds that, while they may be taxed differently, have pretax cash flows that are close substitutes for the par bonds. In this way the par taxables inherit, from a pricing standpoint, tax advantages that they do not actually have. This, of course, creates a difference between the price of par taxables and taxable investors' implicit valuations of their cash flows. Unlike dealers, however, the ability of individual investors to arbitrage this relationship is limited by the need to use investment income to offset investment interest expense. Thus, dealers and institutions in the model dominate the pricing of taxables *relative to each other*, while individuals determine their pricing *relative to tax exempts.*

The basic intuition the model formalizes is quite simple. Investors holding both taxables and municipals might not regard coupon income as fully taxable at the margin because of offsetting investment interest expense elsewhere in their portfolios. These implicit benefits will increase with maturity, pulling down the taxable par yield curve at the long end where more of the present value of a bond is due to the coupon stream. Within the model I develop here, these implicit tax advantages depend in a specific fashion on maturity, but the pricing relations that emerge seem to do a good job of capturing the differences in curvature between the two yield curves. The model also makes fortuitous use of the characteristics of available data, since it prices taxable bonds off the par yield curve for tax exempts.

To evaluate the fit of the model and compare it to the performance of the traditional "tax-adjusted" yields, I use both descriptive techniques and formal estimation with generalized method of moments (GMM). The model produces yields that match the moments of the observed series much better than do the tax-adjusted yields. While the model replicates the differences in curvature between the two yield curves quite well, it continues to understate slightly the long tax-exempt yields (or overstate the taxable yields). This failure is a more severe problem in periods surrounding major changes in the tax law, when investors were anticipating decreases in tax rates.

I also estimate the tax rate implied by the model and evaluate its fit statistically, using fitted zero-coupon rates from the U.S. Treasury market and the corresponding yields for par tax exempts predicted by the model. Since the tax rate appears in these yields in a nonlinear manner, I perform this estimation by using GMM. Given the very limited cross-sectional data available, estimating the model requires assuming that the tax rate of the marginal investor be fixed across time. Because of this, the model is rejected when it is tested over long periods of time when the tax schedule is extremely progressive or includes major changes in statutory rates. In the period subsequent to the 1986 tax reform, however, the model fits quite well. During this period the actual tax regime is reasonably close to that assumed in the model. The marginal rate changes very little over the relevant range of income levels, and there is no distinction between the rates for long-term and short-term capital gains. In all periods the yields predicted by the model match the data much better than do the traditional tax-adjusted yields.

The article is organized as follows. In Section 1, I briefly review the available evidence, discuss its puzzling features, and evaluate other explanations that have been offered in the literature. In Section 2, I then develop the model. In Section 3, I present a simple optimization problem that generates, as a solution, the behavior described in Section 2. In Section 4, I descriptively reevaluate the evidence in light of the model, and, in Section 5, I estimate the implicit tax rate and test the model's fit. In Section 6, I summarize the conclusions.

1. Review of the Puzzle and Available Explanations

We begin by discussing some standard bond-pricing relationships that serve as a basis for the subsequent analysis.

Consider the choices faced by an investor with a fixed tax rate of τ. Let M_T denote the coupon on a par tax exempt maturing at T, and let C_T be the coupon on a par taxable with the same term. Let d_t be the discount factor for after-tax cash flows at date t, or the expected marginal rate of substitution of the investor between today and date t. We also need the following assumptions.

Assumption 1. *There are taxable and tax-exempt bonds priced at par available for each maturity.*

Assumption 2. *All bonds are riskless and noncallable.*

Assumption 3. *Bonds are priced on the basis of cash flows generated if held to maturity. In particular, the options to realize capital losses early are ignored.*

Assumption 4. *Investors trade freely and without frictions in all markets simultaneously, and taxation of long and short positions is completely symmetric so that investors are "marginal" on all bonds.*

Assumption 5. *Over the life of each bond the tax rate and the tax regime will not change.*

Under these circumstances the following equations must hold, if the bonds are to be priced at par:

Table 1
Descriptive statistics for yields of various maturities

		Means				Standard deviations			
Period	Maturity	Muni	Treas	Model	Muni $1-\tau$	Muni	Treas	Model	Muni $1-\tau$
1/50-12/90	1	3.42	5.86	5.86	5.86	1.92	3.36	3.35	3.36
	5	4.04	6.31	6.51	6.92	2.11	3.24	3.31	3.70
	10	4.49	6.42	6.80	7.69	2.32	3.18	3.30	4.08
	20	5.03	6.50	7.04	8.99	2.55	3.18	3.32	4.55
	30[1]	5.23	5.83	7.07	8.99	2.55	3.03	3.27	4.55
1/82-12/85	1	5.89	10.32	10.32	10.32	1.02	1.89	1.89	1.89
	5	7.49	11.62	11.98	13.12	0.98	1.49	1.59	1.88
	10	8.64	11.83	12.65	15.13	1.01	1.34	1.48	2.01
	20	9.64	11.97	13.01	16.88	1.11	1.18	1.50	2.24
	30	9.80	11.89	12.97	17.17	1.14	1.11	1.54	2.31
1/87-12/90	1	5.39	7.66	7.66	7.66	0.79	0.92	0.92	0.92
	5	6.07	8.30	8.28	8.65	0.52	0.64	0.58	0.64
	10	6.59	8.55	8.61	9.42	0.39	0.57	0.44	0.58
	20	7.17	8.72	8.88	10.26	0.43	0.49	0.51	0.81
	30	7.29	8.63	8.84	10.43	0.45	0.51	0.53	0.85

All series are reported monthly and quoted as annual percentages. Municipal and Treasury yields are from Salomon Brothers, *Analytical Record of Yields and Yield Spreads.* Yields in the Model column are artificial taxable yields calculated using the municipal yields, Equations (9) and (13), and the tax rate τ, implicit in the spread between the one-year municipals and Treasuries. Yields in the muni/$(1-\tau)$ column are artificial taxable yields calculated using the traditional equivalent after-tax rule.

[1] Thirty-year Treasury series begins 4/53.

$$1 = \sum_{t=1}^{T} M_T d_t + d_T, \tag{1}$$

$$1 = \sum_{t=1}^{T} C_T(1-\tau) d_t + d_T. \tag{2}$$

Any solution to these two linear equations in two unknowns will have the property that

$$M_T = C_T(1-\tau). \tag{3}$$

Since the yield on a par bond equals its coupon rate, the par yield curve for the tax exempts should be a constant multiple of the yield curve for the par taxables.

This appears to be grossly violated in the data. This is documented in Tables 1-3 and Figure 1. The data employed here, as in virtually all the existing work on tax-exempt bonds, are taken from Salomon Brothers' *Analytical Record of Yields and Yield Spreads* (1991). This source provides monthly estimates of yields for municipals in different default risk categories for various maturities. I consider here only the "Prime" scale, which Salomon Brothers identifies as equivalent to Aaa. The relationships above, which apply to par bonds, should be

Table 2
Descriptive statistics for term premiums at various maturities

Period	Maturity	Means			Minima			Maxima		
		Muni	Treas	Model	Muni	Treas	Model	Muni	Treas	Model
1/50-12/90	5	0.62	0.45	0.65	-0.05	-2.43	-1.53	2.25	2.01	2.92
	10	1.06	0.56	0.94	0.00	-3.08	-2.53	3.30	2.50	3.63
	20	1.60	0.64	1.18	-0.05	-3.28	-2.03	4.75	2.90	3.80
	30[1]	1.80	0.60	1.21	0.00	-3.41	-2.89	5.00	2.94	3.73
1/82-12/85	5	1.60	1.30	1.66	1.05	-0.18	0.67	2.25	2.01	2.92
	10	2.75	1.51	2.33	2.10	-0.30	1.16	3.30	2.50	3.63
	20	3.75	1.65	2.69	2.95	-0.43	1.33	4.75	2.90	3.80
	30	3.91	1.57	2.65	3.05	-0.63	1.22	5.00	2.94	3.73
1/87-12/90	5	0.68	0.63	0.62	-0.05	-0.21	-0.42	1.50	1.65	1.64
	10	1.20	0.88	0.95	0.20	-0.31	-0.44	2.35	2.25	2.38
	20	1.79	1.06	1.21	0.50	-0.39	-0.39	3.10	2.50	2.84
	30	1.91	0.97	1.18	0.60	-0.49	-0.45	3.35	2.41	3.00

All series are reported monthly and quoted as annual percentages. Term premiums are calculated by subtracting from the yield for given month and maturity the corresponding yield for a one-year maturity. Municipal and Treasury yields are from Salomon Brothers, *Analytical Record of Yields and Yield Spreads.* Yields in the Model column are artificial taxable yields calculated using the municipal yields, Equations (9) and (13), and the tax rate τ, implicit in the spread between the one-year municipals and Treasuries.

[1] Thirty-year Treasury series begins 4/53.

Table 3
Mean implicit tax rates

Maturity	Subperiods		
	1/50-12/90[1]	1/82-12/85	1/87-12/90
1	41.0	42.7	29.8
5	36.2	35.4	26.9
10	30.8	26.9	22.8
20	23.2	19.5	18.7
30	18.7	17.6	15.5

The implicit tax rate in month t is calculated as the τ_t solving $M_{Tt} = C_{Tt}(1 - \tau_t)$, where M_{Tt} is the yield on a municipal of maturity T, and C_T is the yield on a taxable of maturity T. All series are reported monthly. Municipal and Treasury yields are from Salomon Brothers, *Analytical Record of Yields and Yield Spreads.*

[1] Thirty-year Treasury series begins 4/53.

closely approximated by this data since, as Salomon Brothers state, "These scales are derived from typical municipal scales for new issues coming out at approximately those dates" (emphasis added). The taxable securities employed are the Treasury yield series reported in the same source. These yields are "read from monthly yield curves which we have prepared. Whenever there is a choice of coupons the curves follow the yields of higher coupon issues in the longer maturities." As discussed here in detail, the broad features of the data appear very robust to these choices.

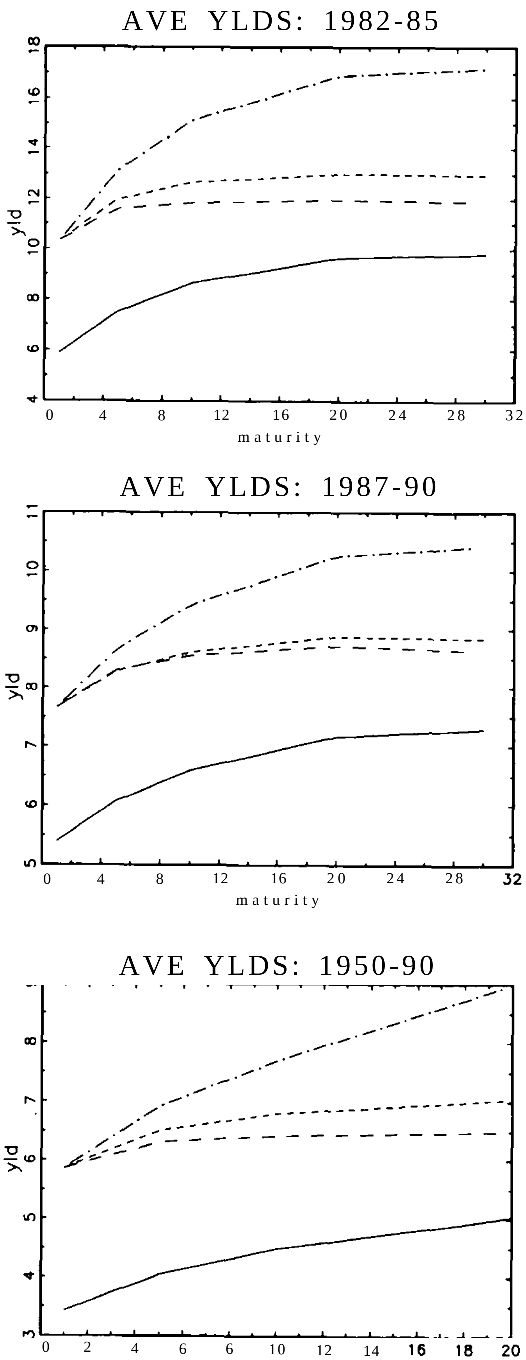
AVE YLDS: 1982-85
yld
maturity
AVE YLDS: 1987-90
yld
maturity
AVE YLDS: 1950-90
yld

In Table 1, I provide descriptive statistics for municipal and Treasury yields in the columns headed “Muni” and “Treas.” Statistics are reported for the entire history of this data, January 1950 to December 1990, and for two recent subperiods.[2] During the most recent of these, 1987-1990, the personal rate for high tax investors of 28 percent was close to the corporate rate of 34 percent, and realized capital gains were all taxed at the personal rate. The other subperiod covers the years 1982-1985. During this time there was considerable discussion of tax reform. The maximum corporate rate was 46 percent, and the maximum personal rate was 50 percent. Realized long-term capital gains were taxed at 40 percent of the ordinary personal rate.

In each of these periods, the average yields for Treasuries increase less dramatically with maturity than do those of the prime municipals, as can also be seen in Figure 1, where average yields are plotted against maturity. The relative steepness of the yield curves is also evident in the columns labeled “Muni” and “Treas” in Table 2, which shows statistics describing the term premiums in the taxable and tax-exempt yield curves. At all maturities the average term premiums are higher for the tax-exempt bonds. Note particularly the minimum values of the term premiums. Although these are substantial negative numbers for the taxable bonds, the municipal yield curve has been flat but never really inverted. In Table 3, I report average values of the implicit tax rates, which solve (3) given the taxable and tax-exempt yields for each maturity and date. These rates are high enough at short maturities to represent a plausible marginal rate for some market participant. At longer maturities the implicit tax rates seem too low to be the marginal rate of any investor who could reasonably be expected to also hold municipals.[3] Note that in all three periods

[2] The remaining columns of Tables 1 and 2, and the associated curves in Figure 1, will be discussed in Section 4.

[3] The effects of state taxes complicate the pricing of municipals and make it more difficult to identify any particular statutory rate as the relevant marginal rate that should determine relative par yields. State taxes segment the municipal bond market because many states only exempt income from bonds issued by their own municipalities. Two factors should be kept in mind, however. First, state taxes should serve to increase the implicit tax rates above the statutory federal rates. Yet what is particularly anomalous is that the tax rates implicit in relative yields on long-term taxables and tax exempts seems too *low* to be plausibly related to the rates of the relevant taxable investors. Second, to the extent I have been able to gather evidence on the subject, differential state taxation appears to effect the levels rather than the curvatures of the yield curves. For example, I have obtained fitted yield curves for January 20,1990, through Miller, Anderson, and Sherrerd that report separately a “California Scale,” a “New Jersey Scale,” and a “Treasury Scale.” The New Jersey Scale

←

Figure 1

Plots of average yield at different maturities for several subperiods

In each panel the bottom curve (solid) depicts municipal yields, the second curve (long dashes) depicts Treasury yields, the third curve (short dashes) depicts yields produced by the model in the article, and the top curve (dots and dashes) depicts yields calculated with the traditional tax adjustment.

the implicit tax rates decline monotonically in maturity. All these behaviors are very evident in other subperiods as well.

Most prior work on the relative pricing of taxable and tax-exempt securities has not focused specifically on the difference in the curvatures of the two yield curves. The work that has considered this question has been largely descriptive, documenting the magnitude and pervasiveness of the problem or asking how the relative spreads respond to changes in the economic environment. A recent and comprehensive summary of the evidence is contained in Jordan and Jordan (1989). Poterba (1989) also provides summary evidence of the responsiveness of the spreads to regulatory and tax changes.

There have been a few attempts to address directly the difference in curvature in the literature. These explanations can be usefully categorized as attempts to relax one of the five assumptions cited. I will consider them in order. My intention in this discussion is not to rule out the possibility that any or all of these explanations could have some merit. Rather, it is to argue that none of them provide an obvious or full explanation for the anomaly and that, taken together, these arguments leave much to be explained.

Yield and returns. The first and most obvious possibility is that the difference in curvature is simply an artifact of the nonlinearities in the yield-to-maturity calculation. Fabozzi and Feldstein (1983, p. 311) state:

> *There is a major drawback in employing the equivalent taxable yield formula to compare the relative investment merits of a taxable and tax-exempt bond. Recall . . . that the yield-to-maturity measure assumes that the entire coupon interest can be reinvested at the computed yield.*

The relationships being considered here, however, are for bonds at (or close to) par value. They do not involve the yield as such; they involve direct comparison of the present value of after-tax cash flows. It is simply fortuitous that, for par bonds, equating these present values involves equating the after-tax coupons, and, for a par bond, the coupon is equal to the yield.

Default and call risk. The arguments that led to Equation (3) require all bonds to be riskless; and U.S. Treasury securities arguably have different default risks than prime-grade municipals.

is simply a translation, by 10 basis points, of the California Scale. For each of them the longest bonds yield 1.3 percent more than the shortest bonds. On the other hand, the term premium for the Treasury Scale is less than 0.5 percent. Both of these arguments suggest that state taxation is unlikely to explain the behaviors that are the focus of this study.

The property that the yield spread for long maturities is much lower than for short maturities, however, is also quite evident in comparisons between municipals and corporates. For example, Trczinka (1982) shows that across all ratings categories the implicit tax rate falls with maturity. More recently, Jordan and Jordan (1989) find that over the period January 1976-May 1987 implicit tax rates for maturities over 25 years averaged 26.3 percent when prime municipals were compared to AAA corporates, and 25.24 percent when medium-grade municipals were compared to A corporates. The same comparisons for lo-year maturities give implicit tax rates of 34.12 percent and 31.69 percent, respectively. Comparing one-year prime municipals to one-year Treasury bills gives a rate of 44.51 percent. In short, the differences across maturity are considerably larger than the differences across quality levels.

Trczinka (1982) argues that changes in the tax and regulatory environment alter the risk of municipals as a group relative to corporates and yet do not occasion wholesale rating changes. It may be, then, that the ratings authorities do not attempt to maintain parity in default risk in their ratings scales for corporates and municipals. Instead, they focus on the default risk of municipals relative to each other.

While this argument has some merit, it is unclear why differences in default risk should have such a pronounced and pervasive effect on the *curvature* of the yield curves. The theoretical work of Merton (1974) on risky zero-coupon bonds, which has been elaborated in Lee (1981) and Pitts and Selby (1983), shows that higher default risk may decrease, rather than increase, the slope of the yield curve in the range of the longer maturities. With coupon bonds the yield curves are smoothed, but the general shape is preserved, as is evident in the numerical results of Kim, Ramaswamy, and Sundaresan (1989). Empirically, Sarig and Warga (1989) find that corporate zero-coupon bonds show relations between yield, default risk, and maturity that are similar to the patterns produced by Merton's model. Yawitz, Maloney, and Ederington (1985) attempt to account explicitly for default risk in a formal econometric model. They find, however, that their estimates of the implicit tax rate still increase with maturity across all ratings categories.

Differences in callability could, in principle, also provide an explanation. Most municipal bond issues are callable, and the yields Salomon Brothers estimates may reflect the prevalence of call features in the market. Since these estimates are for bonds with "new issue" characteristics, the longer bonds would have higher yields because of the call risk. Of course, the same argument could also be made concerning corporate issues, which are typically callable as well and which do not manifest the large term premiums the muncipals display. Further, there have been attempts made in the literature to control for the callability. Kochin and Parks (1988) investigate six specific noncallable, high-grade general obligation issues of states. Since tax-exempt bonds are issued "in series" the tombstone advertisements from the *Wall Street Journal* provide what is, in effect, a par yield curve. On comparing these data to contemporaneous Treasury yields, they find the same strong evidence that implicit tax rates are much lower at long maturities.

Tax options. The foregoing comparisons involve discounted after-tax cash flows if the bonds are held to maturity. Options that accrue to the buyer of a par bond to realize capital losses prior to maturity have values that depend on maturity, which could distort relative yields.

Note that for this argument to have much importance there must be significant asymmetries in the capital gains treatment of taxable and tax-exempt bonds, asymmetries that are to the disadvantage of the tax exempts. Unlike coupon income, capital gains on municipals are taxed. Moreover, it seems unlikely that the option to realize losses early, per se, would tend to favor municipals. First, the lower coupons for the municipals lead to longer durations and more volatility than comparable par taxables. Further, investors buy long municipals in part as a hedge against increases in future tax rates. If tax rates fall, or interest rates rise, so that the price of a municipal purchased at par falls, the investor can realize the loss and roll the proceeds over into a current issue. All subsequent interest then goes untaxed.[4]

There is, though, one such asymmetry in the tax code, referred to as the amortization feature. An investor purchasing a taxable bond at a premium can amortize the premium annually against ordinary income. The premium on a tax exempt cannot be so amortized. Other things being equal, then, long taxables will be more valuable, and have lower yields, because of this feature. Constantinides and Ingersoll (1984) noted that this might explain the narrow yield spreads at long maturities. This possibility is explored in detail in Jordan and Jordan (1989). Their simulations show a steeper term premium for the municipals, but they assume a 50 percent ordinary tax rate and a capital gains rate of zero. Thus, the only tax option available is due to the amortization feature; buy and hold is optimal for the municipals. This regime corresponds to Tax Scenario III in Constantinides and Ingersoll (1984), who refer to their Tax Scenario IV as the "opposite extreme" (p. 334). Under their Scenario IV, where capital gains, losses, and coupon income are all taxed at the same rate, Constan-

[4] Ken Dunn suggested this argument to me.

tinides and Ingersoll argue that the value of a municipal is the same as a taxable bond with the same after-tax coupons. It will not pay to realize a gain through a wash sale today in order to amortize a premium against income in the future when the tax rate on the gain and the future income are the same. This situation corresponds to the post-1986 tax environment, yet the differences in curvature for the taxable and tax-exempt yield curves are still very much in evidence in the post-1986 data.

Clienteles. Prior to the 1984 tax reform, commercial banks were allowed to simultaneously hold short-maturity municipal bonds and deduct interest expense on the deposits that financed these positions. Fama (1977) argued that the implicit tax rates at short maturities should, therefore, reflect the corporate tax rate of the banks. Longer-maturity municipals, held by individuals with lower tax rates, might then have higher yields relative to taxables than the short-term tax exempts.

Obviously, the first problem with this explanation is that it does not apply in the post-1984 environment, even though the high term premium for the municipals persists. Further, as noted in Buser and Hess (1986), the implicit tax rate in short maturities is on average lower than the corporate rate, and it is much more variable. The model outlined here, although it has some attributes of a “clientele” story, explains the curvature with a single tax rate.

Changing tax regimes. A changing tax regime is a risk ignored in the preceding analysis, but, generally speaking, it is a risk shared by holders of taxables. Indeed, it is often argued that municipals are more valuable because they serve as a hedge against changing tax rates. One possibility that obviously would adversely affect municipals is that Congress would simply eliminate the tax-exempt status of municipals. It seems unlikely that this would explain the historical persistence of the differences in curvature, which are evident even in periods where the tax regime has been stable. One must also keep in mind the probability that existing bonds would be “grandfathered” if there was a change in regime.

2. The Model

The discussion in the previous section suggests that as long as investors in high tax brackets are marginal across par municipals and par taxables simultaneously, it is difficult to explain the observed behavior of the yield curves. The model developed here proceeds by relaxing Assumption 4 and, specifically, the assumption that taxation of long and short positions is completely symmetric for individuals. I assume, as the tax code in fact stipulates, that interest deductions from borrowings and short positions cannot exceed investment income from the positions these borrowings support, and I argue that this limits investors' ability to be marginal on all bonds simultaneously. A single representative investor is still at the margin across municipals and some portfolio of taxables, however. Thus, unlike Fama (1977), the model does not rely on different marginal investors at different maturities. The issue raised here is the specific portfolio of taxable bonds the investor would, at the margin, choose to hold. As discussed in the introduction, taxable investors should choose to hold portfolios structured to reduce the impact of taxation on the cash flows received.

Taxable investors who choose to purchase taxable cash flows have a number of means to do so.

1. They can purchase coupon payments, which are fully taxed.
2. They can purchase zeros, with original issue discount, that generate tax liabilities each year as the discount is amortized over the life of the bond.
3. They can purchase coupon bonds.
4. They can buy principal payments with a market discount, which is taxed as a capital gain.

I argue that it is advantageous to do the latter, and then I evaluate the consequences for the relative valuation of municipals and taxables. In fact, investors would prefer to hold the fourth type of cash flow long and short the other three, but the assumption regarding limitations on interest deductions, mentioned earlier, will preclude this. An additional question is why one would confine attention just to these four alternatives. There may be other ways of generating tax-advantaged cash flows besides the strategies described here. Indeed, it is not clear that these strategies are in any sense *optimal*. The next section discusses this problem in more detail.

Consider a position that generates certain payments only on a single date and that is taxed only at that single date. Pretax this is a pure discount bond, but original-issue pure discounts are not actually taxed in this way. They generate tax liabilities each year as the discount is amortized over the life of the bond. A position of the sort we are considering could be created, though, through transactions in coupon bonds of the same maturity.

Suppose there is a 10 percent par bond that matures at date T and a 5 percent bond that trades at a market discount and matures on the same date. Let $B_{5\%}$ denote the price of the discount bond. (The price of the 10 percent is, of course, \$1 by assumption.) Both bonds have \$1 face value. A position that is long two of the 5 percent bonds and short one 10 percent will have no cash flows, and no net tax liability, prior to maturity. Before tax this is just a \$1 pure discount bond. Let P_T denote the cost of such a position. For the example under con- consideration, $P_T = 2B_{5\%} - 1$. After tax the position pays

$$1 - \tau(2 - 2B_{5\%}) = 1 - \tau[1 - (2B_{5\%} - 1)] = 1 - \tau(1 - P_T). \quad (4)$$

Such a position is, in effect, a pure discount bond that is not taxed until maturity.

To a taxable investor willing to hold this position, its price must equal its discounted after-tax value; hence,

$$P_T = [1 - \tau(1 - P_T)]d_T \quad (5)$$

or

$$P_T = \frac{(1 - \tau)d_T}{1 - \tau d_T}. \quad (6)$$

From the standpoint of a taxable investor, payments generated by a strategy such as that just described are clearly tax-advantaged relative to coupon payments arriving at the same date. Compared to cash payments received in the form of maturity payments, which for par bonds are tax-free, they are tax-disadvantaged. From the standpoint of a dealer or a tax-exempt institution, however, cash flow from any of these three sources is equivalent. Any difference in prices between the identical packages of cash flows generated in these different ways creates arbitrage opportunities for them. In particular, consider a par taxable bond and the sequence of pure principal payments, constructed as before, that replicates it on a pretax basis. They should have the same price, so

$$1 = \sum_{t=1}^{T} C_T P_t + P_T, \quad (7)$$

or, using (6),

$$1 = \sum_{t=1}^{T} C_T \frac{(1 - \tau)d_t}{1 - \tau d_t} + \frac{(1 - \tau)d_T}{1 - \tau d_T}. \quad (8)$$

A taxable investor with these discount factors would value a municipal of the same maturity as in (1), which I repeat here for easy reference:

$$1 = \sum_{t=1}^{T} M_T d_t + d_T. \quad (1)$$

Equations (1) and (8) can be used to derive the par taxable yield curve from the par municipal yield curve. With par municipals of every maturity we can solve for the after-tax discount factors via the recursive relation

$$d_T = \frac{1 - M_T(\Sigma_{t=1}^{T-1} d_t)}{1 + M_T}, \tag{9}$$

$$d_1 = \frac{1}{1 + M_1}. \tag{10}$$

The resulting d_t's can be used, in conjunction with Equation (8), to solve for the yield on a par taxable. Calculations of this sort are performed in Section 4.

Now consider how the coupon levels, or yields since these are par bonds, that satisfy (1) and (8) must behave relative to pairs that satisfy Equation (3), $M_T = C_T(1 - \tau)$. The discount factors applicable to pretax payments, derived in (6), reflect the tax shield attributable to the cost of the position. That is, from (6), $P_T > (1 - \tau)\, d_T$. The results of the model simply reflect the fact that a portfolio of these positions, which duplicates the pretax payments of a par taxable bond, will have higher after-tax cash flows on the coupon dates and lower after-tax cash flows at the maturity date. As the maturity of the bond increases, the maturity payments become a less important component of the present value of the bond, which pulls the taxable and tax-exempt yield curves together for longer maturities.

To see this directly, rearrange (8) to give

$$\frac{1 - d_T}{1 - \tau d_T} = C_T(1 - \tau) \sum_{t=1}^{T} \frac{d_t}{1 - \tau d_t}, \tag{11}$$

and rearrange (1) to obtain

$$M_T \sum_{t=1}^{T} d_t = 1 - d_T. \tag{12}$$

Using (12) to substitute for the numerator on the left-hand side of (11) and rearranging gives

$$M_T = C_T(1 - \tau) \frac{\Sigma_{t=1}^{T} [d_t/(1 - \tau d_t)]}{\Sigma_{t=1}^{T} [d_t/(1 - \tau d_T)]}. \tag{13}$$

Note that when $T = 1$ the ratio multiplying the after-tax coupon on the right side of (13) equals 1. Thus, for one-year bonds, the expression $1 - M_T/C_T$ gives the tax rate of the representative investor, and the model is consistent with the standard after-tax equivalent yield calculation.[5] For longer maturities this ratio overstates the marginal

[5] Note that the maturity for which the traditional after-tax yield calculation holds exactly depends on the assumed payment interval for coupons and for taxes. I am implicitly assuming here, and in the subsequent empirical work, that these coincide and that the payment interval is one year. This assumption seems natural in terms of the structure of the tax code and simplifies the calibration exercises in Section 4. To match the semiannual payment interval for the bonds would require assuming that taxes are paid twice a year. This feature of the model is unavoidable, since the value

tax rate. The two sums in the ratio on the right side of (13) differ only in the subscript on the discount factor in the dominators of the terms being summed. As long as forward rates are positive, $d_T < d_t$ for $t < T$, and $d_t/(1 - d_t) > d_t/(1 - d_T)$. The numerator thus exceeds the denominator. The longer the maturity, the more terms accumulate in the numerator that are large relative to the terms in the denominator, making M_T increasingly large relative to $C_T(1 - \tau)$.

In frictionless markets, aside from tax considerations, dollars can be purchased for delivery at date T in a manner that shields part of the payment from taxation. Prices will adjust so that, at equilibrium, taxable investors will be marginal, or indifferent, between tax-exempt dollars at date T and taxable dollars from these positions. On a pretax basis, though, the principal-only payments created as suggested can be used to replicate coupon-bearing bonds, and it is assumed that arbitrage by dealers forces prices of the two pretax streams together.[6] On the other hand, coupon-bearing taxables with long maturities, held in isolation, will look unattractive to taxable investors at the prices produced by the model.

The question then is, why would taxable investors not short par taxables and invest in portfolios of synthesized zeros at market discounts that replicate the par bonds on a pretax basis? Note this is not pure tax arbitrage if the maturity is finite. The replicating portfolio will involve a tax liability at maturity, whereas the par bond's maturity payment is return of principal, which is tax-free. If, however, discount factors are monotonically declining in maturity, so that forward rates are positive, the present value of the tax disadvantage at maturity will eventually become negligible relative to the tax savings. In a model with no market frictions, taxable investors would wish to trade in this way to the point that (2) and (8) were equated. At that point, however, arbitrage opportunities would be available to dealers and tax-exempt institutions, who would trade to equate pretax values.

There are at least two reasons why the activities of dealers should dominate in this process.

1. It is surely the case that dealers have a comparative cost advantage in these transactions.
2. More importantly, the tax code limits the trading individuals can do that would drive (2) and (8) into equality. These trades involve

of the deferrals created through these strategies depends on when payments are received and taxes paid within the year. Of course, the consequences are small. on the order of the differences between annual and semiannual compounding. The qualitative features of the analysis, specifically the maturity dependence of the ratio in (13), will hold regardless of how time is indexed.

[6] Dealers can and do arbitrage in this way across Treasury securities. That is, after all, how zero-coupon bonds were first marketed. The extent to which such arbitrage succeeds in driving prices into line is ultimately an empirical issue. The anomalous pricing of coupon versus principal strips suggests that some barriers to this sort of arbitrage exist.

selling par coupon bonds short and buying the synthetic zeros that replicate the par bonds on a pretax basis. The interest deductions generated in creating the synthetic zeros exactly match the interest income from the position. The coupons are "stripped." If such a position is combined, however, with additional short positions in a coupon bond, the interest deductions generated by the short positions would exceed the interest income from the long positions. The tax code requires that such interest deductions be limited to interest income from the securities that the borrowing is incurred to purchase or carry.[7] This limits the ability of individual investors to trade in a direction that, if their positions could be infinite, would force equality between (2) and (8). Note that unlike the appeal to higher "market frictions" for individuals, this limitation to their ability to trade against dealers and institutions does not interfere with their ability to construct the tax-preferred principal-only portfolios in the first place.

Some remarks are in order concerning the realism of the assumptions employed to obtain these pricing results. The model assumes frictionless short selling and bond markets sufficiently complete as to have multiple maturities available at any date. Obviously, therefore, the actual construction of positions that are taxed as in (4) would be problematic. This, in turn, calls the pricing relation in (8) into question. At least two factors merit consideration in this regard, however.

- Many models in financial economics require freedom to short-sell with full use of the proceeds and unrealistically complete markets, notably the mean-variance CAPM, the arbitrage pricing theory, and the Black-Scholes option pricing model. Despite the evident impossibility of implementing the transactions these models envisage, such as holding the "market" portfolio, the models have proved helpful in explaining observed behaviors of asset prices, identifying certain behaviors as anomalous, and aiding and evaluating financial decision making. When these models fail to explain important features of the data, the frictions they assume away are a natural place to look for an explanation, but the presence of such frictions need not invalidate a model a priori.
- The arguments should be robust to the following extent. They appeal to the intuitive and general notion that investors attempt to shield coupon income from taxation by generating offsetting interest expense elsewhere in their portfolios. Portfolios that achieve this are attractive to taxable investors, and this should affect how taxables are priced relative to tax exempts. In particular, such opportunities make coupon-bearing instruments relatively more attractive, and as the

[7] See Fabozzi (1987, pp. 48-49) for a discussion of these Issues.

maturity of the bond lengthens the coupon stream becomes an ever larger portion of the present value of the after-tax payment stream. There should, therefore, be a consequent maturity effect of the sort observed. This might not be reflected in the pricing of taxables relative to each other, simply because when the same pretax cash flows have different prices dealers have arbitrage opportunities. They are relatively unhindered by the factors, including both transactions costs and limits in the deductibility of investment interest expense, that constrain the ability of individual investors to exploit such an opportunity.

3. Optimal Taxable Positions

The previous section *assumes* taxable agents are at the margin between holdings of par municipals and portfolios of taxable bonds composed of synthesized zero-coupon positions at market discounts. It then describes the consequences of this for the relative valuation of par taxables and par municipals. These positions are obviously tax-advantaged relative to simply holding par taxables, but they may not be optimal. If optimal portfolios have very different characteristics, the theory will be misleading. Of course, other tax-motivated portfolio strategies are also likely to involve attempts to substitute other, tax-advantaged cash flows for fully taxable coupon income. They should, therefore, have similar implications for relative yields on longer-maturity par bonds. Nevertheless, it may be informative, and perhaps reassuring, to delineate a set of circumstances under which the behavior described in the previous section–stripping coupons to transform ordinary income into capital gains–can be associated with the solution to an optimization problem.

This section provides an example of such circumstances. This analysis is not offered as a comprehensive, definitive model of bond market equilibrium. The solutions to multiperiod investment consumption programs under differential taxation will depend on agents' time and risk preferences, as well as on the range of bonds available and the assumed tax regime. Any simple tax-based strategy, such as that exploited above, can only be optimal in an extremely stylized environment. Understanding the characteristics of this environment, however, can be helpful in evaluating the reasonableness of the behaviors assumed and their pricing consequences, and it is in this spirit that the following example is developed.

Assume that there are three dates, $t = 0,1,2$, and investors have three bonds available, which are priced at a discount, at par, and at a premium. Let B_a be the price of the discount bond with coupon C_a. The par bond has a price of unity and coupon C_b. The premium bond has price B_c and coupon C_c. As previously, I assume here that the relative prices of the three bonds are free of tax effects because of dealer arbitrage; hence,

$$B_a = C_aP_1 + (1 + C_a)P_2, \tag{14}$$

$$1 = C_bP_1 + (1 + C_b)P_2, \tag{15}$$

$$B_c = C_cP_1 + (1 + C_c)P_2 = 1 + \alpha(B_a - 1), \tag{16}$$

where P_t is the pretax discount factor and

$$\alpha = \frac{C_c - C_b}{C_a - C_b}. \tag{17}$$

The last equality follows from the fact that the premium bond can, on a pretax basis, be replicated as a portfolio of the other two and must be priced accordingly.

We will ask what is the cheapest way to generate a given payment at date $t = 2$ subject to the constraint, first, that the payment at $t = 1$ be nonnegative and, second, that the deduction generated by premium amortization must be employed to offset interest income from the portfolio itself. Let $y = \{y_a, y_b, y_c\}$ denote the positions taken in the three bonds, and use (16) to substitute for the price of the premium bond in all expressions. We must solve the cost minimization problem,

$$\operatorname*{Min}_{\{y_a, y_b, y_c\}} \; y_aB_a + y_b + y_c[1 + \alpha(B_a - 1)], \tag{18}$$

subject to

$$\begin{aligned} 1 \leq\; & y_a[1 + C_a(1 - \tau) - \tau(1 - B_a)] + y_b[1 + C_b(1 - \tau)] \\ & + y_c[1 + C_c(1 - \tau) + (\tau/2)\alpha(B_a - 1)], \end{aligned} \tag{19}$$

$$\begin{aligned} 0 \leq\; & y_aC_a(1 - \tau) + y_bC_b(1 - \tau) \\ & + y_c[C_c(1 - \tau) + (\tau/2)\alpha(B_a - 1)], \end{aligned} \tag{20}$$

$$0 \leq y_aC_a + y_bC_b + y_c[C_c - \tfrac{1}{2}\alpha(B_a - 1)]. \tag{21}$$

This problem assumes that the premium on bond 3 can be amortized linearly. The objective function is just the cost of the portfolio. The first constraint, (19). states the portfolio must make a unit after-tax payment at maturity; the second, (20), that it make a nonnegative payment at $t = 1$. The premium amortization is included in both these payments. The third constraint, (21), requires some explanation. Without such a constraint, this problem would offer, at zero cost, tax deductions at $t = 1$ financed by capital gains at $t = 2$. A portfolio that holds a units of bond a and $1 - \alpha$ units of bond b pays a coupon of C_c, a maturity payment of \$1, and generates a capital loss of $\alpha\ (B_a - 1)$ at $t = 2$. (Note that $\alpha < 0$.) Shorting this portfolio and buying the premium bond generates a capital gain, from the position in the discount bond, of $\alpha\ (B_a - 1) > 0$, at $t = 2$, but this is offset by amortization of the premium, which in total is of equal magnitude. Since the premium is amortized over two periods, while the market discount is only recognized as a gain at $t = 2$, this zero-cost position generates deductions at $t = 1$ in exchange for taxes at $t = 2$. An agent with positive time value for after-tax income will always engage in such trades.

The constraint (21) is intended to proxy for the features of the tax code that limit such deductions to investment income. If the investor had such income from other sources, the right-hand side of the constraint would be some negative number, and there would be positive holdings of the premium bond. This constraint would bind, however, and the investor's *marginal* valuation of the bonds would be consistent with that suggested by the discussion to come. (Formally, the dual variables in the subsequent problem do not depend on the right-hand side of this constraint.) Alternatively, with a zero position in the premium bond, the constraint can be interpreted as requiring that interest deductions must be used to offset interest income from elsewhere in the portfolio.

The following proposition shows that solutions to this problem involve the creation of synthetic zero-coupon bonds that are priced at a market discount. An agent solving such a problem, as part of a general consumption-investment program, would value dollars from taxable bonds relative to par tax exempts as suggested in the previous section.

Proposition 1. *If forward rates are positive, then the solution to the cost minimization problem involves zero holdings of the premium bond and coupon stripping—a long position in the discount bond and a short position in the par bond in proportion to their relative coupons.*

Proof. Let the prices of the three bonds be denoted as the vector d, and let the 3×3 matrix of coefficients in the constraints be denoted A. The cost minimization problem can be written

(primal)
$$\min_{y} d'y \tag{22}$$

subject to

$$Ay \geq e_1, \tag{23}$$

where e_1 is a 3-vector with unity as its first entry and zeros elsewhere. This problem has a dual of

(dual)
$$\operatorname*{Max}_{\rho} e_1'\rho \tag{24}$$

subject to

$$A'\rho = d, \tag{25}$$

$$\rho \geq 0. \tag{26}$$

A solution to the equality constraint in the dual, r^*, will satisfy

$$\rho^* = (A')^{-1}d. \tag{27}$$

Since this equation uniquely determines r^* it will represent a solution to the dual as long as it satisfies $r^* \geq 0$. Substitution into the equality constraint and algebraic simplification show that the values

$$\rho_1^* = \frac{P_2}{\tau P_2 + (1 - \tau)}, \tag{28}$$

$$\rho_2^* = \frac{P_1}{\tau P_1 + (1 - \tau)}, \tag{29}$$

$$\rho_3^* = \tau P_1(\rho_2 - \rho_1) \tag{30}$$

satisfy the required system of equations. The values of ρ_1^* and ρ_2^* are clearly positive. Cross-multiplication within the difference on the right-hand side of the expression for ρ_3^* shows that it is positive as long as $P_1 > P_2$ which is implied by positive forward rates. So ρ^* is feasible for the dual and must be the unique solution.

Consider $y^* = A^{-1}e_1$ as a solution to the primal. It is clearly feasible, since it satisfies the constraints with equality. To show that it is the solution, we must show that the objective functions for the dual and the primal are equal at the proposed solution. But

$$d'y^* = d'A^{-1}e_1 = \rho^{*\prime}e_1 = e_1'\rho^*. \tag{31}$$

It remains to show that y^* has the stated characteristics. The solution to the constraints as equalities is unique. The values

$$y_a^* = \frac{1}{1 - \tau(1 - B_a) - C_a/C_b},$$
$$y_b^* = -(C_a/C_b)y_a^*,$$
$$y_c^* = 0 \tag{32}$$

clearly solve the constraints as equalities. They must, therefore, represent the solution. By inspection y^* has the characteristics claimed for it in Proposition 1. ■

4. Comparisons with Historical Yield Curves

The model derived in Section 2 can be used to price taxable bonds given the par yield curve for tax exempts. The focus on par yields is fortuitous, since the available data for tax exempts comes in this form. In this section, I compare the yield curves thus generated to yield curves for U.S. Treasuries, using as input the par yield curves for prime-grade municipals provided by Salomon Brothers (1991).

From Equation (9) it is evident that, given a par tax-exempt rate for each maturity, we can solve for the after-tax discount factors for each date. Then, given a tax rate τ, (8) is a linear equation in only one unknown, C_T, the yield on a par taxable. By solving this equation for C_T, we can compare the par taxable yields produced by the model to the Treasury data.

Two things are needed for this exercise that are not at hand. First, for the tax rate of the "marginal" investor I simply take the tax rate implicit in the one-year Treasury rate relative to the one-year municipal rate. From Equation (13), with $T = 1$, it is evident that the model predicts that the one-year implicit tax rate is equal to the marginal rate. An alternative would be to choose the tax rate to minimize some statistical metric across maturities. This approach is pursued in the next section. The simpler procedure here amounts to putting all the weight on a single maturity at each date, so any noise idiosyncratic to that maturity will be passed through the entire yield curve via the tax rate. This choice seems a reasonable place to start, however, for descriptive purposes.

The second problem is that the Salomon Brothers data only provide yields for selected maturities.[8] To use Equations (8) and (9) to derive the taxable yield curve in terms of the tax exempts, par municipal rates for each maturity are needed. In the next section I use fitted yield curves for Treasury pure discounts to overcome this problem, but this limits the range of maturities and time periods that can be considered. For the moment, therefore, I adopt the simpler strategy of just interpolating the tax-exempt yield curve linearly and investigating the model's "fit" informally through descriptive statistics and graphs.

Note that while both of these choices are naive, their naivete would tend to work against the model. The first will overweight one partic-

[8] From 1950 to 1979 these maturities are 1, 2. 5, 10, 20, and 30 years. From 1980 to 1988, yields for 15-year bonds are included. In 1989 a maturity of three years Is substituted for that of two years, and the 15-year maturity is dropped.

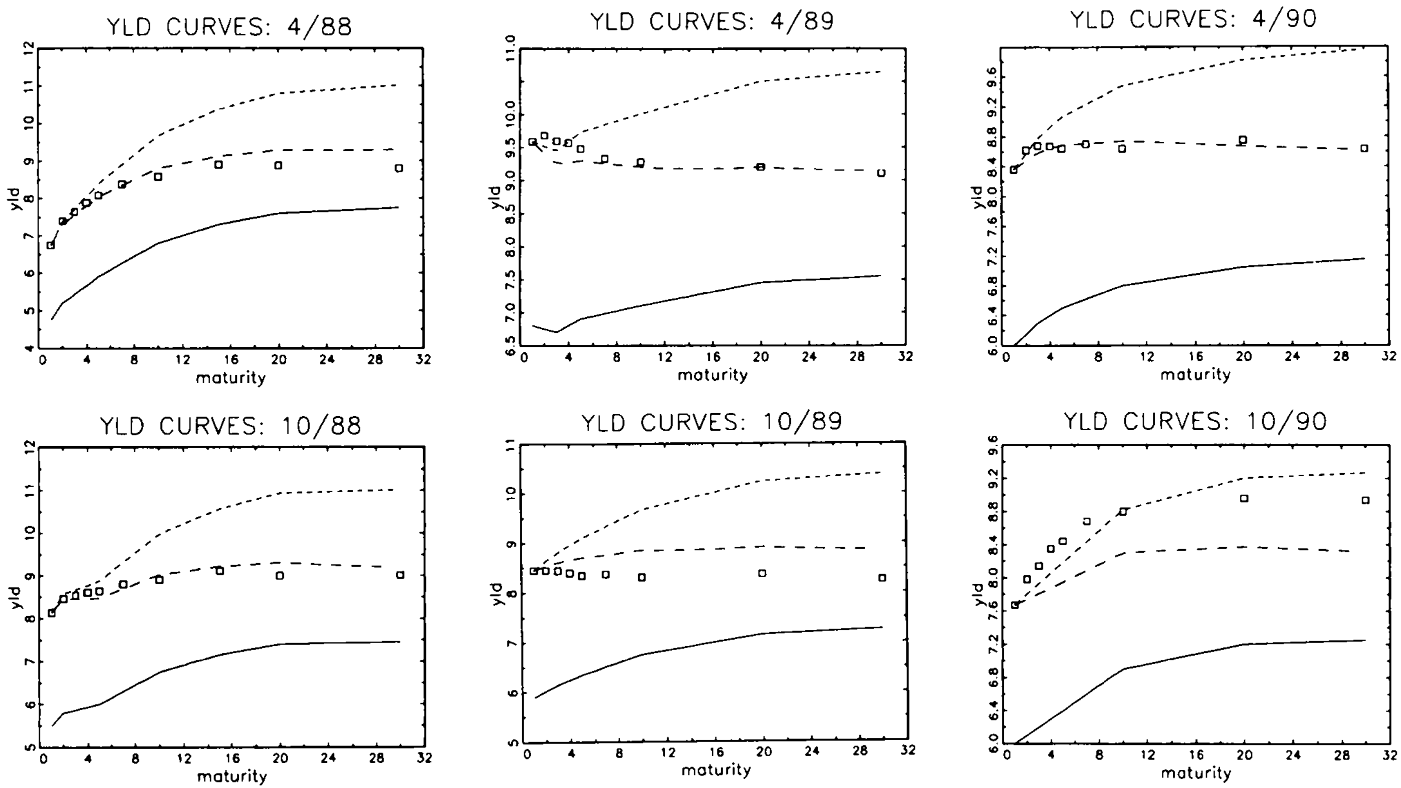

Figure 2
Plots of yield curves at semiannual intervals
In each panel the bottom curve (solid) depicts municipal yields, the second curve (long dashes) depicts yields produced by the model in the article, and the top curve (short dashes) depicts yields calculated with the traditional tax adjustment. Treasury yields are represented as boxes. Implicit tax rates used in calculations are 29.63 percent (4/88), 32.35 percent (10/88), 29.09 percent (4/89), 30.18 percent (10/89), 28.23 percent (4/90), 21.77 percent (10/90).

ular yield relative to the others. The second will straighten out the municipal yield curve for intermediate maturities.

Figure 2 shows plots of the yield curves calculated as described previously at quarterly intervals for the years 1988-1990, which is the historical period when the tax code most closely approximates the very simple one envisioned by the model. By historical standards, there is relatively little progressivity in the tax schedule during this period. Further, aside from rates for very low incomes, the statutory rates are relatively close to each other, within the range of 28 percent to 34 percent. Other complications discussed in Section 1, such as arbitrage by commercial banks and tax options, are also likely to have less of a contaminating effect during this time period. Thus, if the model has any descriptive power, one would expect it to perform well here.

The panels of the figure plot, at semiannual intervals over the last three years of data, the municipal yield curve, linearly interpolated, the municipal curve divided by 1 minus the tax rate, the yield curve produced by the model, and the Salomon Brothers Treasury yields. The model reproduces the curvature of the Treasury rates quite well. When it misses, it overstates more often than understates the actual taxable rates, but the predicted and actual taxable yields are generally quite close. The discrepancies between the two do not appear to be a failure to pick up some aspect of the way the taxable and tax-exempt yields fluctuate relative to each other. In rising term structures the two taxable yield curves both rise less steeply than the municipal yields do, and the inversion in early 1989 is effectively reproduced by the model. At the end of 1990, unreasonably low tax rates implicit in the short-term yields lead the model to understate the taxable yields at long maturities. During this period, anxiety about the credit risk of several major municipalities, including New York and Philadelphia, became acute. These concerns centered on their ability to roll over their short-term paper and appear to have increased short-term municipal yields more dramatically than long-term yields.

These features are evident at other times as well. There are a few periods when the model systematically overstates the long-term taxable yields, 1982 and 1986, for example. The obvious explanations are the proposed tax cut in 1982 and the tax reform legislation being considered in 1986. On the other hand, the extreme inversions of the yield curve in 1980 and 1981 are very accurately replicated by the model. During the late 1950s the model tends to fit less well, but it misses on both the high and low sides. These problems are largely attributable to the volatility of the implicit tax rate. Interest rates were very stable and very low during this time, and revisions to the data appear to have been somewhat infrequent and incomplete. As a ratio, however, the implicit tax rate is very sensitive to the noise thus introduced.[9]

Tables 1 and 2, in Section 1, show descriptive statistics calculated using monthly data for artificial yields produced by the model that are analogous to the statistics describing the municipal and Treasury yields. These are reported in the columns headed "Model." The means of these artificial yields are also plotted in Figure 1. For comparison purposes, Table 1 and Figure 1 also report statistics for artificial yields calculated by using the traditional tax adjustment achieved by grossing the municipal yields up by $1 - \tau$. In the period 1987-1990, as can be seen in Table 1 and Figure 1, the model on average only slightly overstates the taxable yield. The period 1982-1985, on the other hand, is one in which the model did not fit very well. The average yields and term premiums are off by much more, and the discrepancies are clearly larger for longer maturities. At this time, of course, there was both considerable uncertainty about future tax rates and expectations that they would fall over time. The effects of these anticipated tax reforms are evident in the reported minimal values for the term premiums in Table 2. The minima over the overall 30-year period testify to the model's ability to replicate inversions for the taxable term structure. In 1982, however, part of the inversion of the taxable yield curve appears to have been due to anticipation of tax cuts, and the model, because of the steep slope of the contemporaneous tax-exempt yield curve, produces term premiums that are positive. For the overall period, the means are matched on average much better than in 1982-1985, but it is still evident that the average differences increase with maturity. Across all periods, the model produces taxable yields that are much closer to the actual ones than those obtained by simply multiplying the municipal yields by $(1 - \tau)^{-1}$.

The close match achieved in terms of curvature suggests that the model would accurately reproduce differences in volatility across maturities, and this is also borne out in Table 1. The long municipal rates are quite volatile. Simply dividing by $(1 - \tau)$ scales this volatility up too much in comparison to the behavior of the long Treasuries through time. The nonlinear transformation associated with the model smooths out much of this volatility and leads to time-series behavior that is much closer to what is observed.

5. Estimation and Tests

This section evaluates how closely the model fits the historical data using a statistical metric. I estimate the "implicit tax rate" of the

[9] This noise is evidenced by the fact that during this period the implicit tax rate is occasionally negative.

marginal investor, using both the model described in Section 2 and the traditional approach of equating after-tax yields. Since the model expresses par tax-exempt yields as a nonlinear function of the tax rate, Hansen's generalized method of moments (GMM) is employed in this estimation.

The objective is to compare statistically the yields on tax-exempt bonds produced by the model to the yields reported in the Salomon Brothers data. Calculating the par yield curves using the model requires, at a given point in calendar time, either the pretax discount factors $\{P_{ts}\}_t$ or the after-tax discount factors $\{d_{ts}\}_t$. The subscript s denotes calendar time, while t refers to maturity. I begin with the pretax discount factors and use the fitted zero-coupon rates created by Bliss and Fama, which are available on the CRSP government bond tapes. This database provides discount factors for maturities of one to five years at the end of each month from June 1952 through December 1990. For a given tax rate, we can find the price of a tax-exempt dollar at each maturity by inverting Equation (6) to obtain

$$d_{ts}(\tau) = \frac{P_{ts}}{1 - \tau(1 - P_{ts})} . \tag{33}$$

These prices can then be used in Equation (1) to calculate the coupon rate on a par tax-exempt bond, which matures T periods from date s, as a function of the tax parameter:

$$M_{Ts}(\tau) = \frac{1 - d_{Ts}(\tau)}{\Sigma_{t=1}^{T} d_{ts}(\tau)} . \tag{34}$$

Note that substitution of (33) into the above gives, for $T = 1$, the tax-exempt rate as the return on the taxable bond with price P_{1s} times $(1 - \tau)$.

Let M_{Ts} denote the observed yield at date s on a T-period par tax exempt. The model does not speak to the sources of errors from the theoretical yields, so the choice of moment conditions to use in estimation is somewhat arbitrary. I simply assume, first, that the unconditional expectation of the errors is zero at the "true" tax rate τ_0:

$$E\{M_{Ts} - M_{Ts}(\tau_0)\} = 0. \tag{35}$$

Second, to account for the possibility that the yields are nonstationary, I assume the expectation of the error given past yields is zero and use the following moment condition:

$$E\left\{\frac{M_{Ts} - M_{Ts}(\tau_0)}{M_{T_{s-1}}}\right\} = 0. \tag{36}$$

The Salomon Brothers data provide yields for prime-grade municipals with maturities of one, two, and five years,[10] along with selected longer maturities. Given that fitted zero rates are only available with maturities up to five years, I can calculate $M_{T_s}(\tau)$ for maturities of one to five years. Thus, using these data sources, I have orthogonality conditions of the form (35) or (36) for three maturities, T = 1,2,5, that can be employed. With only one parameter to estimate, however, the system is overidentified with these three moment conditions.

The estimation proceeds by choosing the tax rate to minimize a quadratic form in the vector of sample analogs to the three orthogonality conditions (35) or (36). Here these sample moments are simply the three sample means, calculated in time series, of the differences between the predicted yields at a given tax rate and the observed yields, or the sample means of these differences normalized by the previous observed yield. I employ the standard two-stage procedure. In the first stage a consistent estimate of the tax rate is obtained by using the identity matrix to form the quadratic form to be minimized. In the second stage the weighting matrix was calculated from the Newey-West procedure. The results reported in the tables all employ three lags in forming the weighting matrix. All these estimations were also performed with a wide range of different lag lengths, with no qualitative consequences for the results.[11]

Table 4 shows the results of estimating this model over the entire period for which the two data sources overlap and various subperiods during which the tax regime was more or less fixed.[12] The table shows estimates using both levels [Equation (35)] and ratios [Equation (36)] because, like most interest rates series, these yields exhibit behavior suggestive of a unit root. The method of moments estimation requires stationarity and ergodicity of the series involved. Using the entire sample, estimated first-order autocorrelations for each maturity exceed 98 percent, though the sample autocorrelations appear to decline

[10] In 1989 and 1990 yields are provided for maturities of one, three, and five years.

[11] The estimations of the model using (35) and using (36) were performed with lag lengths of 0, 1, 3. 4, 6, 12, and 24 months using the entire sample period and for the post-1986 period. Over the other subperiods lag lengths of 3, 4, and 6 were employed. For the traditional model, lag lengths of 0, 3, and 6 months were used. Using more lags in all cases increased the standard errors of the parameters and reduced the x^2 from the specification test, but in no case were these differences large enough to affect the conclusions. At all lag lengths the model is rejected at standard significance levels in every period, except for the post-1986 period. The traditional model Is rejected In all periods and produces uniformly higher x^2-statistics. Even at the highest lag lengths employed, the rejections reported at three lags in the tables persisted at least at the 5 percent level. Point estimates of the tax rate were very robust to changes in the number of lags. For example, estimates obtained using (36) ranged from 40.201 to 40.907 percent when the whole sample was employed, and the lag length ranged from 0 to 24. Using the post-1986 data the tax rate ranged between 31.997 and 31.043 percent.

[12] These subperiods are the times during which the statutory rates were fixed and the maximal rates could plausibly have been the marginal rate.

Table 4
GMM estimates of tax rates and specification tests using the model developed in the article

Time period	Moment conditions $E\{M_{Ts} - M_{Ts}(\tau)\} = 0$			Moment conditions $\frac{E\{M_{Ts} - M_{Ts}(\tau)\}}{M_{T_s-1}} = 0$		
	n	Tax rate	$x^2(2)$	n	Tax rate	$x^2(2)$
6/52-12/90	462	40.373%	48.832	461	40.201%	40.9%
		(13.543)	(0.000)		(10.698)	(0.000)
1/87-12/90	48	32.065	0.606	47	31.762	2.669
		(5.188)	(0.739)		(4.899)	(0.263)
1/65-12/80	192	44.013	25.274	191	44.006	26.072
		(11.003)	(0.000)		(9.852)	(0.000)
1/82-12/85	48	43.745	10.852	47	44.397	11.244
		(3.396)	(0.004)		(3.633)	(0.004)

Standard errors and *P* values are in parentheses. The weighting matrix was estimated using the Newey-West procedure with three lags. The orthogonality conditions are imposed for maturities of one, two, and five years, leaving two degrees of freedom for the x^2. The *n* column reports the number of monthly observations employed in the estimation.

geometrically with the lag.[13] Such high autocorrelations may be due to a unit root or to regime shifts associated with changing tax policy. Over subperiods of more stable tax regimes the autocorrelations are lower and decline much more quickly.[14] Univariate Dickey-Fuller tests using the Newey-West correction produce only very weak evidence against a unit root in these series. The F-statistics, depending on the lag length employed, are generally close to the critical value for 10 percent and rarely exceed the 5 percent critical value.[15] Attempts to estimate the model in differenced form were unsuccessful, leading frequently to negative tax rates and extremely high standard errors. Evidently, most of the information about the tax rate is eliminated in differencing. The ratios in the moment condition (37), however, show no evidence of the persistence the yields themselves display. The sample autocorrelations of M_{Ts}/M_{Ts-1}, T = 1,2,5, are negative at the third and fourth lags and less than 60 percent at the first lag. Estimating the model using this orthogonality condition, therefore, provides a check on whether the results based on the levels are robust concerning the issue of stationarity. Table 4 is reassuring in this regard. Parameter estimates and x^2-statistics are very similar across the two sets of moment conditions. Notice that the estimated tax rates all

[13] For example, the one-year rates have sample autocorrelations at higher lags that are uniformly lower than the first-order autocorrelation (.9818) raised to the lag. Nevertheless, at the 24th lag, the sample autocorrelation is still .639.

[14] For instance, in the post-1986 period the sample autocorrelations for the one-year yield are .892, .243, and -.368 at lags of 1, 12, and 24 months.

[15] For example, for maturities of one, two, and live years. the F-statistics are 5.434, 6.104, and 5.563 with a three-lag correction. They are 6.245, 6.337, and 5.012 with a 48-lag correction. The critical values, as reported in Fuller (1976), are 5.34 (10 percent), 6.25 (5 percent), and 7.16 (2.5 percent).

seem plausible, given the subperiod employed. To the extent that the asymptotic standard errors are reliable, the tax rates also appear to be estimated with reasonable precision.

The x^2-statistics in Table 4 provide a goodness-of-fit test of the model. The overidentifying restrictions are rejected at high confidence levels over the entire period. This is hardly surprising. Over these years there were several major revisions in tax law, including dramatic changes in the statutory rates. The model attempts to fit the data with a single tax rate that is constant through time. Thus, the model cannot be expected to fit even as well as suggested by the descriptive procedures in Section 4, which allowed the tax rate to vary period by period. Similarly, over the subperiod January 1965 to December 1980 the maximum statutory rate was 70 percent. It seems reasonable to expect that the rate of the "marginal investor" would lie well below this level and, given the progressivity of the schedule below the maximum rate, that it should vary through time. Thus, it is asking a great deal of the model to fit a single tax rate over this whole period to the data, and the model is, in fact, rejected.

Over the two more recent periods the model performs much better. In the most recent period, subsequent to the 1986 tax reform, the model is not rejected at conventional levels of significance. During this period the marginal rates were constant over a much larger range of incomes at the higher end of the spectrum. The tax schedule coincides reasonably well to that which is envisioned in the model. Over the relevant range of income the marginal rate is relatively constant. The estimated rate is within one standard error of either (i) the rate of 28 percent that applies to high income levels or (ii) that rate plus the 5 percent surcharge that applies to taxpayers in the 33 percent "bubble." Further, a number of confounding factors present in earlier years have been eliminated from the current tax regime. These include the distinction between long- and short-term capital gains rates, and the opportunities for banks to arbitrage over the short-maturity ranges. I argued in Section 1 that these effects were insufficient to explain the magnitude or persistence of the anomaly and, in particular, its presence in recent data. The absence of such confounding effects in the current regime, however, may explain the model's relative success in the most recent periods.

Unfortunately, the most recent period also provides us with relatively few time-series observations. To evaluate whether the consequent lack of power is the primary reason for the failure to reject in recent periods, I also estimated the traditional equivalent after-tax yield model, using the same methods and sample periods. The results are presented in Table 5. It is evident that the fit of this model, as measured by the magnitude of the x^2-statistic, is worse in every period.

Table 5
GMM estimates of tax rates and specification tests using the traditional model that equates after-tax yields

Time period	Moment conditions $E\{M_{T_s} - C_{T_s}(1-\tau)\} = 0$ n	Tax rate	$\chi^2(2)$	Moment conditions $\frac{E\{M_{T_s} - C_{T_s}(1-\tau)\}}{M_{T_s-1}} = 0$ n	Tax rate	$\chi^2(2)$
6/52-12/90	462	40.167% (10.877)	89.402 (0.000)	461	40.123% (9.730)	98.105 (0.000)
1/87-12/90	48	33.263 (4.263)	13.368 (0.001)	47	33.661 (4.439)	13.283 (0.001)
1/65-12/80	192	37.611 (10.073)	47.937 (0.000)	191	41.696 (9.198)	52.901 (0.000)
1/82-12/85	48	4 1.573 (3.447)	14.943 (0.001)	47	42.335 (3.626)	15.399 (0.000)

Standard errors and P values arc in parentheses. The weighting matrix was estimated using the Newey-West procedure with three lags. The orthogonality conditions arc imposed for maturities of one, two, and five years, leaving two degrees of freedom for the x^2. The n column reports the number of months of data employed in the estimations.

In particular, during the post-1986 period the traditional model is rejected at conventional significance levels, which suggests that the model's success over this period is not solely attributable to a lack of power to reject.

6. Conclusion

I examined the anomalous behavior of the taxable and tax-exempt yield curves. Long municipal yields appear to be too high relative to the equivalent after-tax yield that can be earned in Treasury or corporate bonds. After considering various explanations of the problem, I proposed a simple model that relates the yields of taxable bonds to the yields on par tax exempts. The model implies that the ratio of the tax-exempt yield to the taxable yield decreases with maturity, which is consistent with observed phenomena, such as inverted yield curves for taxables and contemporaneous rising yield curves for tax exempts. Statistical and descriptive comparisons between the yields predicted by the model and observed yields show that the model has some success in explaining the apparent anomalies in the relative curvatures of the two yield curves.

References

Buser, S. A., and P. J. Hess, 1986, "Empirical Determinants of the Relative Yields on Taxable and Tax-Exempt Securities," *Journal of Financial Economics,* 17, 335-355.

Constantinides, G., and J. Ingersoll, 1984, "Optimal Bond Trading with Personal Taxes," *Journal of Financial Economics,* 13, 57-64.

Fabozzi, F. J., 1987, "Federal Income Tax Treatment of Fixed Income Securities," in F. Fabozzi and T.M. Pollack (eds.), *Handbook of Fixed Income Securities,* Dow Jones-Irwin, Homewood, Ill.

Fabozzi, F., and S. Feldstein, 1983, "Municipal Bonds," in F. Fabozzi and T.M. Pollack (eds.), *Handbook of Fixed Income Securities,* Dow Jones-Irwin, Homewood, Ill.

Fama, E., 1977, "A Pricing Model for the Municipal Bond Market," working paper, University of Chicago.

Fuller, W. A.. 1976, *Introduction to Statistical Time Series,* Wiley, New York.

Jordan, B. D., and S. D. Jordan, 1989, "Tax-Timing Options and the Relative Yields on Municipal and Taxable Bonds," working paper, University of Missouri-Columbia.

Kim, I. J., K. Ramaswamy, and S. Sundaresan, 1989, "The Valuation of Corporate Fixed Income Securities," working paper, Columbia University.

Kochin, L. A., and R. W. Parks, 1988, "Was the Tax-Exempt Bond Market Inefficient or Were Expected Tax Rates Negative?" *Journal of Finance,* 43, 913-931.

Lee, C. J., 1981, "The Pricing of Corporate Debt: A Note," *Journal of Finance,* 36, 1187-1189.

Merton, R. C., 1974, "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance,* 29, 449-470.

Pitts, C. G. C., and M. J. P. Selby, 1983, "The Pricing of Corporate Debt: A Further Note," *Journal of Finance,* 37, 1311-1313.

Poterba, J. M., 1989, "Tax Reform and the Market for Tax-Exempt Debt," *Regional Science and Urban Economics,* 19, 537-562.

Salomon Brothers, 1991, *Analytical Record of Yields and Yield Spreads,* New York.

Sarig, O., and A. Warga, 1989, "Some Empirical Estimates of the Risk Structure of Interest Rates." *Journal of Finance,* 44, 1351-1360.

Trczinka, C., 1982, "The Pricing of Tax-Exempt Bonds and the Miller Hypothesis," *Journal of Finance*, 37, 907-923.

Yawitx, J. B., K. J. Maloney, and L. H. Ederington, 1985, "Taxes. Default Risk, and Yield Spreads," *Journal of Finance,* 40, 1127-1140.