Why are tax-exempt bonds issued at a premium?*

Mattia Landoni
SMU Cox School of Business
lando@smu.edu

January 1, 2015

Abstract

I derive the optimal timing strategy for an investor taking tax gains and losses on tax-exempt bonds (up to a 7% gain over buy-and-hold). I then derive the issuer's optimal coupon rate, which maximizes the tax benefit for a tax-timing investor (up to a 3.5% gain over issuing at par, potentially more than the cost of issuance itself). All these gains are transfers from the U.S. Treasury to local issuers or investors.

The optimal issuance policy is consistent with three unexplained stylized facts: the frequent issuance of premium bonds; "sticky" coupons that don't fall when yields fall; and higher issue prices for longer-maturity noncallable bonds.

*PRELIMINARY—PLEASE DO NOT CITE WITHOUT THE AUTHOR'S PERMISSION. COMMENTS ARE WELCOME.

1 Introduction

The U.S. municipal bond market is the world's largest market for tax-exempt assets.[1] However, few of the municipal bond market's many quirks have been examined in the finance literature.[2] One well-established stylized fact is that many municipal bonds are issued at a premium, i.e. with a coupon rate significantly larger than the yield. By comparison, corporate bonds, government bonds, and even taxable municipal bonds are issued at par very often (Figure 1), i.e. issuers appear to have an incentive to set the coupon near the bond's required yield.

Yet another fact is that coupons are "sticky": as bond yields have fallen to the recent record lows, coupon rates on new issues have remained high, resulting in more premium bonds than usual (Figure 2). To the best of my knowledge, neither one of these facts has been previously considered, let alone explained, in the academic finance literature.

Existing literature on optimal security design focuses mostly on the tradeoff between alleviating agency problems and informational asymmetries on the one hand, and avoiding inefficient liquidation on the other (Anderson and Sundaresan, 1996; DeMarzo and Fishman, 2007; Bhamra et al., 2010): committing to a higher coupon rate is a virtuous signal, but it also puts pressure on the firm's cash flow and it can cause inefficient liquidation. In an empirical paper, Amiram et al. (2014) also mention behavioral and institutional frictions – for instance, the fact that par bonds are "simpler" to understand, hence preferred by retail investors.

None of these theories offers an immediate explanation for the apparent puzzle of premium bonds. Existing models often feature a simplified setup in which the coupon is a very important variable in the issuer's capital structure: often the issuer has only *one* bond outstanding; sometimes this bond is a perpetuity and the coupon rate is truly a proxy for the choice of a whole capital structure. However, for a municipal issuer, the coupon is largely an arbitrary accounting number. Municipal issuers typically issue "series" of perhaps 10-15 or more bonds, each bond with a different maturity and coupon rate. These series are often designed to achieve "level debt service", i.e. equal payments over time – akin to a household taking a fixed rate home mortgage. These level cash flows are then packaged into bonds, perhaps to maximize their value to potential buyers. Thus, the coupon rate of any individual

[1]With a total capitalization of roughly four trillion dollars, the U.S. municipal bond market constitutes about 8 percent of the net wealth of U.S. households and nonprofits, and about 2 percent of the world's financial assets. Total net wealth of U.S. households is about $50 trillion, according to the Federal Reserve Flow of Funds statistics. An estimate by McKinsey Global Institute places the stock of all debt and equity across the globe at $212 trillion in 2010.

[2]A recent notable exception is Ang et al.'s (2010) examination of the "de minimis" effect, i.e. the sudden jump in yields as bond prices fall below an artificial boundary created by the applicable tax rules.

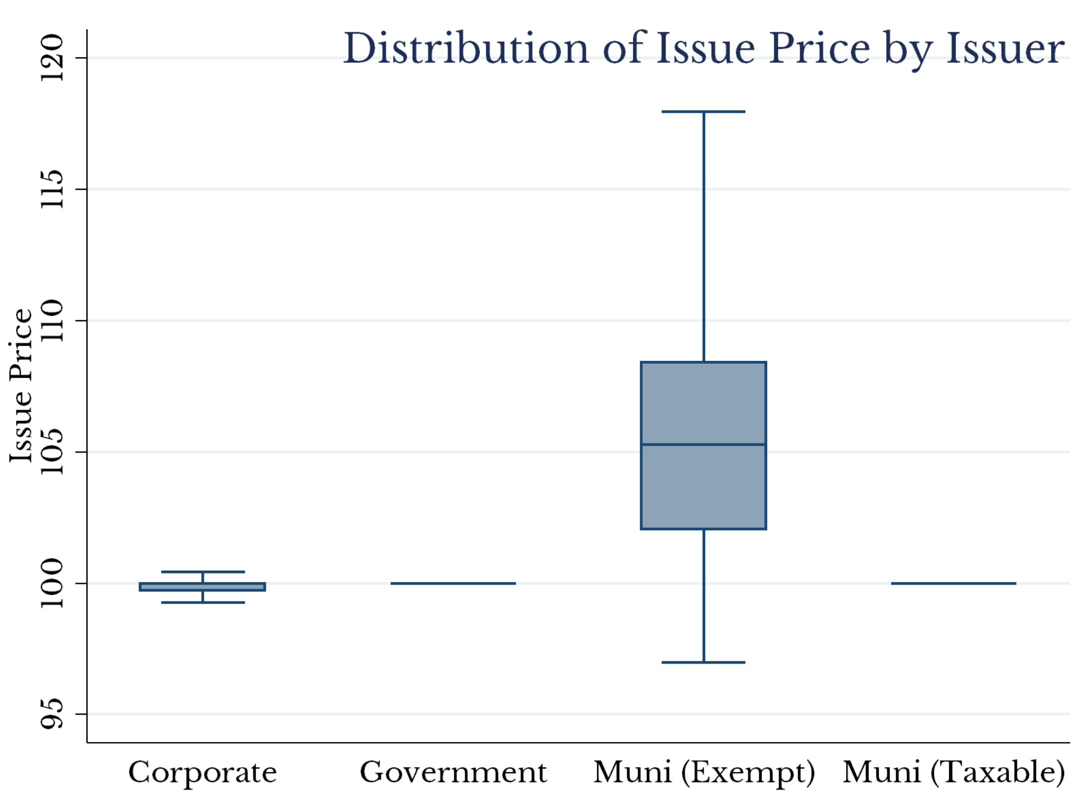

Bond Type	Discount	De Minimis	Par	Premium	N
Corporate	1.54%	49.98%	47.85%	0.63%	15,964
Government	0.66%	17.56%	80.41%	1.37%	11,354
Muni (Exempt)	0%	0.48%	10.02%	89.5%	98,143
Muni (Taxable)	0%	6.01%	74.77%	19.21%	8,280

Figure 1: Tax-exempt municipal bonds are very likely to be issued at a premium (i.e., with a coupon in excess of the required yield), while other types of bonds (including taxable municipal bonds) are typically issued at par.
Such a phenomenon is consistent with issuers maximizing the ex-ante value of investors' option to realize capital gains and losses. The box represents the middle 50% of the empirical distribution of issue price in a sample of over 130,000 bonds held by property and casualty insurance companies between 2004 and 2012.

municipal bond bears little relationship with the actual cash flows of the issuer, depriving the classical agency/information arguments of much of their relevance.

A candidate theory of coupon rates should also accommodate the fact that tax-exempt municipal bonds are often issued at a premium, while taxable municipal bonds are not, and to do so it must leverage some difference between these two types of municipal bonds. In this paper I examine a promising candidate factor that (a) is accounting-related, and thus could influence the choice of arbitrary accounting quantities, and (b) differs across taxable and tax-exempt bonds: taxes.

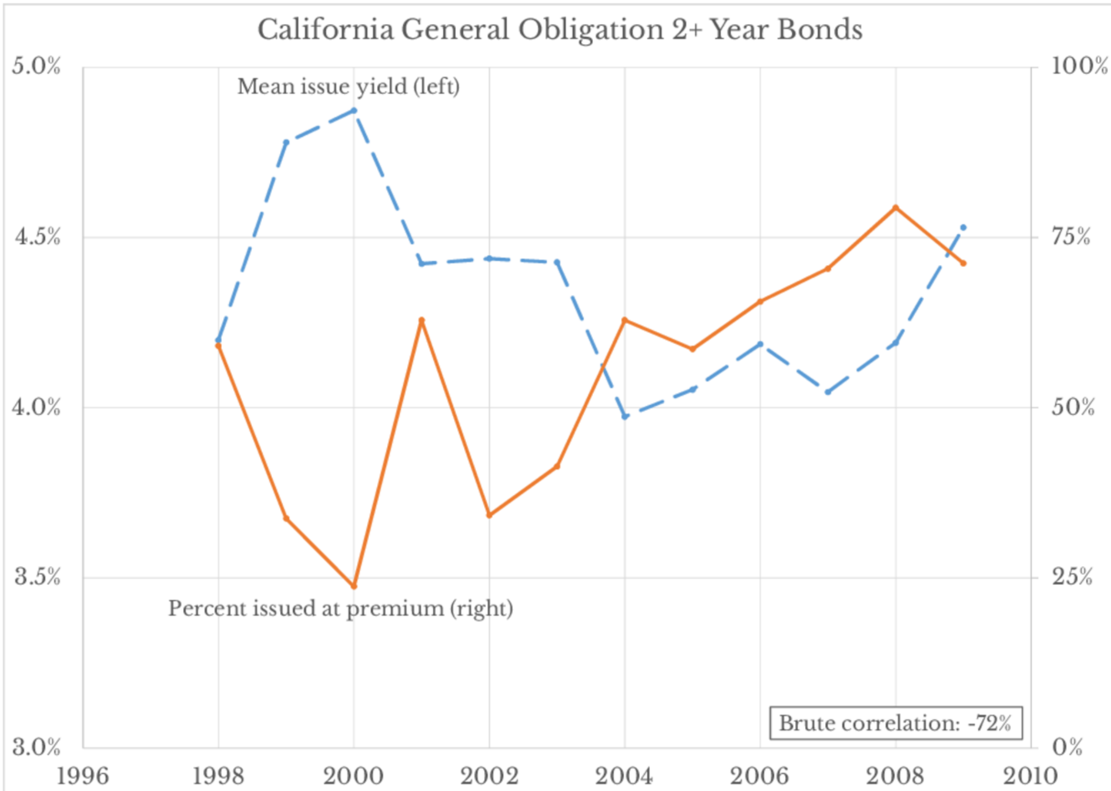

Figure 2: "Sticky coupons": The percentage of bonds issued at a premium is inversely related to the level of issue yields.

To examine the potential role of taxes, I derive rules for the tax-efficient trading and issuance of tax-exempt bonds. Perhaps it is hard to imagine that tax-exempt bonds could be rendered more tax-efficient than they already are; however, capital gains and losses are taxable like they are for every other asset, giving rise to large and distortionary tax trading incentives. In particular, realizing gains is almost always very costly, and realizing losses is very profitable when certain conditions are met.

The first contribution of this paper is to find the investor's optimal trading strategy numerically. Given the current state of the interest rate process, the bond's coupon rate and remaining time to maturity, and the investor's tax rates, liquidity shocks, and transaction costs, I derive when it is optimal to realize gains or losses, and when it is optimal to hold on to the bond.

This paper is not the first to derive a numerical solution for the optimal trading strategy.[3]

[3] Both Constantinides and Ingersoll (1984) and recent work by Kalotay (2013) and Kalotay and Howard (2014) do the same. The magnitude of the tax option value is already noted (and independently verified) by Kalotay and Howard (2014), who use a slightly different methodology. However, Constantinides and Ingersoll focus only on bonds issued at par, and Kalotay and Howard focus on the pricing consequences.

Constantinides and Ingersoll (1984) consider the optimal trading of (both taxable and tax-exempt) bonds with personal taxes. As a way to simplify their task, they choose to focus exclusively on bonds issued at par. This is an entirely innocuous decision with regard to their main goal: obtaining an unbiased measurement of the required return on taxable default-free bonds. Since the price of a bond includes the capitalized value of future tax trading (i.e. the "tax option"), the raw yield is a biased estimate of the required return on the bond.[4]

Constantinides and Ingersoll conclude that, while for taxable bonds the value of the tax option can be large, this value is negligible for tax-exempt bonds. However, the value of trading optimally is far from negligible for an investor in tax-exempt bonds. Simple review of the U.S. tax code suggests that focusing on bonds issued at par needlessly constrains the tax trading opportunities available to the investor; and in reality, tax-exempt bonds are often issued at a premium. For tax-exempt bonds *issued optimally*, the tax option value can easily be worth several percentage points of the issue price. Applying Constantinides and Ingersoll's original purpose to tax-exempt bonds, ignoring the tax option creates a substantial bias in the measurement of the tax-exempt yield curve.[5][6]

The existence of this large tax trading potential is what enables the second and main contribution of this paper: characterizing the tax-efficient coupon rate chosen by the issuer to maximize the value of the investor's tax options, and showing that this rate is high and "sticky", consistent with both real-world "quirks" mentioned above. To the best of my knowledge, this is the first paper showing that security design can be driven by the investor's trading strategy, as opposed to strategic interactions between lenders and borrowers (as in Anderson and Sundaresan, 1996) or other reasons.

The mechanism that drives this result is simple. Under the U.S. federal tax code, for bonds issued at or above par, it is the coupon, not the yield, that is tax-exempt: if a bond is bought in the secondary market at a yield exceeding the coupon, the excess is called "market discount yield" and constitutes taxable income. Clearly, issuing bonds with a large coupon minimizes the probability that later the yield will exceed the coupon. By definition, a bond issued with

[4]In the asset pricing literature, "tax option" generally indicates the ex-ante value of optimally harvesting future gains or losses for tax purposes, versus following a baseline buy-and-hold strategy. Because the investor has the option to realize gains and losses only when it's profitable to do so, tax trading has an option-like payoff; hence, the name.

[5]Interestingly, the bias is an underestimate of the true yield; correcting it would further deepen the "muni puzzle", i.e. the fact that long-term tax-exempt yields appear "too high" compared to their taxable counterparts.

[6]Constantinides and Ingersoll also argue that "no simple characterization of the optimal trading policy is possible;" the optimal strategy is indeed not simple to characterize, as it involves realizing either losses or gains, depending on capital gain and income tax rates, transaction costs, time to maturity, and the prevailing interest rate. However, in Section 3, I show a visual characterization of the optimal trading strategy that provides some intuition.

a large coupon (in excess of the yield required at the time of issue) is a premium bond. Thus, only premium bonds can be tax-efficient. Moreover, the optimal coupon rate is "sticky" because of mean reversion in the interest rate process. When the interest rate is low, the coupon rate should not fall as low, in order to protect the investor from a future expected rise in the interest rate.

The issuer's tax incentive to issue at a premium is large. The value of issuing a bond with the "right" coupon can be up to 3.5 percentage points of the issue price, compared with issuing at par. This magnitude is likely to exceed the cost of issuance itself in many cases; if this benefit is fully captured by the issuer, it has the potential to turn issuance into a profitable business (at the expense of lower revenues for the U.S. Treasury).[7]

Optimal issuance also yields a novel prediction, which I show to be supported by the data: noncallable bond issue prices are increasing with the maturity of the bond.[8] Once again, the mechanism is simple. Long-term bonds have more volatile prices, and a larger issue price is needed in order to accommodate this greater volatility, so that the price will not dip below par (i.e., the yield will not exceed the coupon). Taken all together, these predictions and the corresponding empirical facts constitute strong evidence that taxes are capitalized in asset prices when the marginal investor is taxable.[9]

The paper is structured as follows. In Section 2, I describe and solve numerically a simple model of optimal trading and issuance. The results are described in Sections 3 and 4. In Section 3, I describe the optimal trading strategy for a taxable investor, quantify its value compared to a buy-and-hold benchmark, and provide a simple approximate trading strategy that does not deviate too much from the exact optimal strategy. In Section 4, I show the model's predictions for optimal issuance, and argue that tax-efficient bonds have a high and "sticky", i.e. they stay relatively high when yields fall, and viceversa. I also verify the one novel prediction made by the model: noncallable bond issue prices are increasing in maturity. The

[7] It is also possible that part of the tax benefit is captured by the investors who initially bid for the bonds, if these are not bidding against one another in a fully competitive way. Even so, if any part of the benefit is captured by the issuer, the issuer has an incentive to issue at a premium. For simplicity, throughout the paper, I assume that the bond price is the fully competitive one and all the benefit is captured by the issuer.

[8] This paper focuses on noncallable bonds for simplicity. Empirically, callable bonds are typically issued at a premium as well, but the typical issue price is not strictly increasing in the maturity of the bonds. Enough work would be needed to tease a prediction for callable bonds out of this model, that I leave this task for future research.

[9] Several authors in the past have attempted to produce similar evidence in the context of the U.S. bond market, but have encountered several difficulties. It is hard to find evidence that taxes affect the prices of U.S. Treasury bonds, perhaps because of the prevalence of foreign, sovereign and otherwise tax-insensitive investors (Litzenberger and Rolfo, 1984; Jordan and Jordan, 1991; Green and Ødegaard, 1997; Elton and Green, 1998). It is also hard to reconcile the relative levels of taxable and tax-exempt yields, for a myriad reasons including investor clienteles, differing liquidity, and security features such as call options (for an incomplete and impressionistic list, see Trzcinka, 1982; Green, 1993; Chalmers, 1998; Longstaff, 2011; Jordan, 2012).

coincidence of model and data is striking evidence that even in a market as illiquid and opaque as the municipal market, prices transmit information in a remarkably efficient way.

2 A dynamic model of issuance and trading

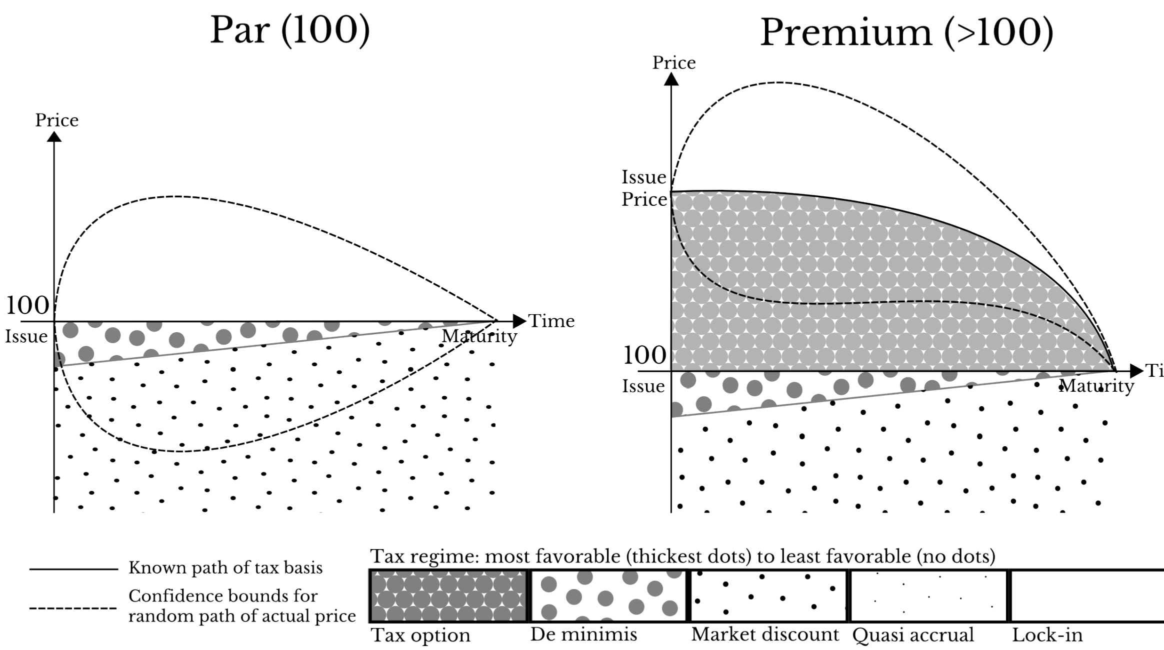

Figure 3: Tax-exempt bonds issued at a premium are endowed with a tax option that tax-exempt bonds issued at par do not have.
(See text for picture interpretation).

In this Section, I identify and describe an "inconsistency" of the tax code that gives rise to all the theoretical and empirical findings described in this paper. I then define and solve numerically a dynamic bond pricing model in order to quantify the consequences of this "inconsistency" for trading and issuance behavior.

2.1 The tax code

First, consider the following example. To simplify the exposition, assume that all investors in the market have the same discount rates and face the same tax rates. An investor buys a tax-exempt bond at issuance with a coupon of 5%. The required yield is also 5%, so that the investor pays a price of 100 (i.e., par value). Consider the following three scenarios.

No Sale Every year until maturity the investor receives a coupon of 5. The investor posts tax-exempt income of 5 and pays no tax. Then, the bond matures and the investor receives 100, the face value. This last cash flow is a return of capital and therefore it is also untaxed. No taxes are paid, and the government's revenue is zero.

Sale at a Gain Suppose that the investor wants to sell the bond and there are 10 years left to maturity. Bond yields have decreased since the time of issuance, and the price of the bond has risen to 110 (a "premium" of 10 over face value). The investor realizes a capital gain of 10. If the capital gains tax rate is 20%, the investor faces a tax liability of 2. The buyer receives a coupon payment of 5 every year, but posts tax-exempt income of 4 only, because she is amortizing the premium paid (10) over the remaining bond life (10 years), and thus $5 - 10/10 = 4$. Upon maturity the buyer receives once again the bond's face value of 100. The buyer thus never pays any taxes, *nor* can she claim a capital loss of 10 to offset the seller's capital gain. Thus, the only taxable event is the gain realized by the seller, and the government's revenue is +2.

Sale at a Loss Finally, suppose that the bond price has fallen to 90 and the investor sells at a loss of 10. Assume for simplicity that there are no limits on the realization of taxable losses, so the seller receives an immediate tax rebate of 2 (10 times the capital gains tax rate of 20%). The buyer holds on to the bond until maturity, then receives 100, the face value of the bond, posting a gain of 10. This gain is considered "market discount income" and taxed at ordinary tax rates. If the buyer's tax rate is 40%, the tax liability (10 years later) is 4. The present value change in government revenue is +0.16.[10]

Thus, if a sale never happens, no taxes are ever paid; when a sale happens and a gain is realized, the government's revenue increases by 2. Thus, realizing gains is very costly. Finally, when a sale happens and a *loss* is realized, government revenue still increases in present value, but only by 0.16. Thus, realizing losses is almost neutral from a tax standpoint.

Second, reconsider the same example, but this time the bond is issued at a premium (say, with an issue price of 120). That is the bond's initial book value for the investor. Slowly, as time goes by, the investor's book value declines due to premium amortization. Ten year before maturity the investor is considering to sell, and the bond's book value is now 110.

In this second example, the first two cases are unvaried. If no sale happens, no taxes are ever paid. If the price of the bond is, say, 120, just as before the buyer realizes a gain of 10, there

[10]For the price of a 10-year, 5% coupon bond to go from 100 to 90, interest rates must rise from 5% to 6.36%. Thus, $-2 + 4(1.0636)^{-10} = 0.16$.

are no consequences for the buyer, and government revenue goes up by 2. The big difference is for the loss case. If the price of the bond is 100, the seller realizes a loss of 10, and gets a rebate of 2, like before. The buyer pays 100 now, and receives 100 when the bond matures, and thus never posts any market discount income and never pays any taxes. Government revenue goes *down* by 2: realizing losses for bonds issued at a premium is very profitable!

Figure 3 depicts the difference between the two examples. In each chart, a point represents a transaction with a time (horizontal axis) and a price (vertical axis). Thicker polka dots represent a greater tax incentive to sell.

The solid black line is the predetermined time path of the tax book value (or “basis”) for an investor who bought the bond at the issue price, and receives 100 upon maturity. The tax basis starts at the issue price at time 0 (a point on the vertical axis), and gradually converges to 100 (a point on the horizontal axis) by the time of maturity. The dashed lines drawing an “avocado” around the tax basis are confidence bounds for the random path of the *actual* market price; unlike the basis, the market price is not known in advance.

Between the examples above and Figure 3, it should be clear that issuing premium bonds is optimal. The larger the issue premium, the larger the “tax option” area in the top-right chart of Figure 3, and the higher the likelihood that future transactions will take place within this area, i.e. under the most favorable tax regime available.[11]

A more intuitive way to explain this outcome is the following: tax law is written in such a way that the coupon, not the yield, is tax-exempt. If, after issuance, a bond trades at a yield higher than the coupon rate, the excess yield is taxable. Issuing with a large coupon prevents this undesirable outcome.

2.2 Preview of less-than-obvious findings

Other implications of the model that will be developed in this section are less obvious. First, as exemplified in Figure 4, often there is a unique, well-defined optimal issue price (i.e., coupon). We know that issuers want “large” coupons to reduce the probability that in the future the bond trades at a yield higher than the coupon. If so, then, why are bonds not issued with an infinite coupon? Why would a given coupon rate ever be “too high”?

An infinite coupon is no problem in practice: one could issue a bond with a face value of $1

[11] This section only considers bonds issued at par or premium. For bonds originally issued at a discount (OID), tax law is even more complicated, but in practice they behave like bonds issued at par. Hence, I ignore OID bonds for simplicity’s sake.

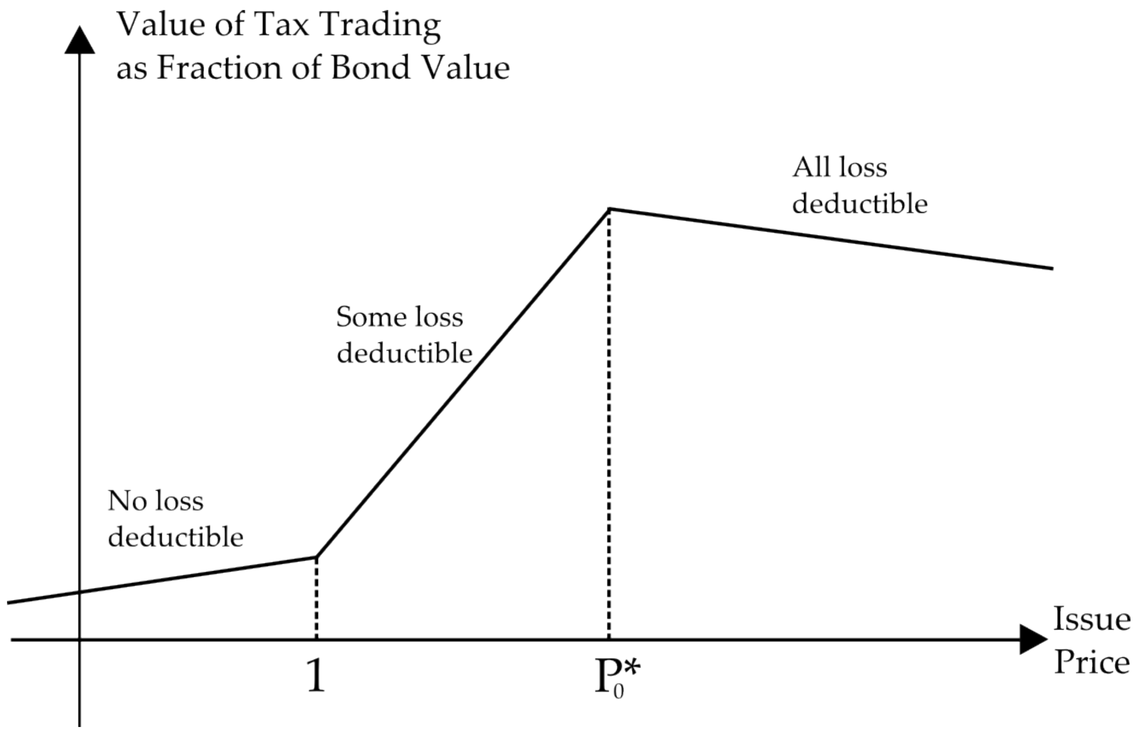

Figure 4: Existence of a finite optimal coupon.
When the issue price is high enough, practically all potential losses are deductible at no future cost. Increasing the coupon further does not add any tax benefits, but it does reduce the bond's duration, dampening volatility, and reducing the tax option value as a percentage of the bond's total value.

that pays a $10,000 coupon (a 1,000,000% coupon rate). However, a bond with a lower coupon has longer duration, hence it has larger gains when interest rates go down, and larger losses when interest rates go up. This is a desirable feature because from a tax trading perspective *losses are a good thing, and gains are indifferent*: the investor optimally refuses to realize gains when gains are present, but always realizes losses when losses are present. This is just an instance of the well-known fact that option holders like volatility; a large coupon shortens the bond's duration, dampening price volatility and ultimately hurting the value of the investor's tax option.

Second, I find that the optimal issue price (as a fraction of face value) should be increasing in the maturity of the bonds and the direction of future interest rate movements. To say that the yield should never be larger than the coupon is the same as to say that the bond should never become a discount bond, i.e. the price should never dip below par. Longer-maturity bond prices are more volatile, so the issue price should be higher to accommodate this extra volatility.[12] Similarly, when yields are expected to rise, the issue price (and the coupon) should be higher to accommodate the higher expected depreciation.

2.3 Model overview

The model described in this section permits to price bonds of arbitrary maturity when taxable investors trade optimally (i.e. realize gains and losses to maximize their after-tax expected return) and bid competitively (i.e. do not have market power, and therefore pay the full value of the bonds they buy). I assume that the issuer and the investor agree on a discount rate, which is a Markov process with 21 states for the one-period interest rate. I model transaction costs and a "realistic" tax code (described below). Given the coupon, c, and the time to maturity, T, I first use backward induction to derive the investor's optimal trading strategy and the exact market price of the bond at any time t prior to maturity as a function of the current state of the interest rate process, r_t.

I then proceed to derive the optimal issuance strategy for the issuer. The market price of a bond at time t, $P_t\left(r_t, T-t, c\right)$, is also the issue price of a bond with maturity $M = T - t$. A fraction W of this price is paid by the investor in anticipation of future tax costs and benefits from realizing gains and losses.[13] I define the optimal coupon c^* as the coupon that maximizes the "subsidy ratio" W, i.e. the subsidy per dollar of principal issued from the federal government to the issuer.

This section describes the model and explains more in detail how it is solved. The next two sections examine the results. In Section 3 I describe the optimal trading strategy for the investor. I also provide a simple approximate sell signal and show that it performs very well. In Section 4, I describe the optimal issuance strategy for the issuer and compare it with summary data on actual issuance.

2.4 The interest rate process

Similar to Constantinides and Ingersoll (1984), I assume that the long-run distribution of the interest rate is uniform. In particular, the interest rate r assumes one of $N_r = 21$ discrete values ($r \in \{0\%,\ 0.5\%,\ \ldots r_{max} = 10\%\}$). Time is divided into discrete one-year ticks. Each year, the

[12] Note that a higher issue price is not the same as a higher coupon rate. For instance, a 5% coupon with a 3% yield will result in a higher price for a 10-year bond compared to a 5-year bond, because the same excess yield (5% - 3% = 2%) is paid to the investor for a longer time.

[13] W can and will sometimes be negative, but can still be maximized.

interest rate has a positive probability to stay at the same level, or to jump up or down by up to two percentage points (e.g., if $r_t = 5\%$, $r_{t+1} \in \{3\%,\ 3.5\%,\ \ldots,\ 6.5\%,\ 7\%\}$. The probability of switching is described by the following table.[14]

Change	-2%	$-1^1/2\%$	-1%	$-^1/2\%$	0	$+^1/2\%$	$+1\%$	$+1^1/2\%$	$+2\%$
Pr (Change)	.05	.07	.11	.15	.24	.15	.11	.07	.05

This simple calibration is meant to roughly mimic the dynamic of the Fed Funds rate without requiring a lengthy explanation. A more sophisticated calibration would, for instance, try to mimic the short-term tax-exempt rate that actually prevails on the U.S. municipal market. Moreover, a yield curve model that yields on average an upward sloping curve would produce more realistic results. However, I do not believe that adding any of these features would change the result substantially.

2.5 Security features, trading, taxes, transaction costs, and investor types

I make several assumptions for simplicity. First, bonds are assumed to be noncallable and pay a coupon only once a year. While most municipal bonds are in fact callable, the presence of a call option makes things more complicated, but does not change the main intuition. For instance, premium bonds are still optimal, and the premium is still increasing in price volatility.[15]

Second, investors can trade the bonds once a year, immediately after the coupon has been paid. As in the simple model of Section D, there is an exogenous probability λ that the investor must trade; however, if realizing gains or losses is optimal, the investor will trade with certainty.

Third, the tax code is a simplified but "realistic" version of the actual U.S. federal income tax code, and I argue that it essentially captures the economics of the actual rules. Market discount income is explicitly taxed as ordinary income. Ordinary income, capital gains/losses, and tax-exempt income are taxed respectively at tax rates τ, τ_G, and τ_E.[16] More details and simplifying assumptions are discussed in detail in Appendix B for the interested reader.

[14] When the current rate is less than two percentage points away from the reflecting boundaries of 0% and 10%, an adjustment is needed because some events in the table are impossible. For instance, if the interest rate is 8.5%, a +2% change would bring the interest rate to 10.5%, which is higher than the upper bound. In this case, the vector is simply truncated and rescaled to sum to one. This adjustment leaves the long-run distribution uniform.

[15] Brown (2011) suggests that the call option itself may be advantageous from a tax standpoint. However, both volatility and the tax basis are likely to be affected by the time to first call, as opposed to the time to maturity. In the end, everything depends on the specific details of the call feature. Understanding the potentially complicated interaction between optimal coupon and optimal call features is left to future research.

[16] Note that τ_E can be nonzero for certain investors, as shown in Table 1.

Finally, trading may or may not be cost-free. Depending on several attributes of the sale, the seller, and the bond, the transaction cost K could be as high as a whole percentage point of the bond's market value. On the other hand, often transaction costs may not matter at all: if an investor is deciding which of two bonds to sell, incurring a transaction cost is inevitable, and it is not part of the decision. Similarly, if the bid-ask spread is due mainly to inventory risk and asymmetric information, the cost of trading for tax purposes (e.g. selling a bond to realize a gain and buying it back immediately) could be very low, because tax trading does not create any of these burdens for the dealer.

Investor Type	τ	τ_G	τ_E	λ	K
Individuals					
Benchmark	40%	20%	0%	0%	0.0%
Frequent Trading, Low Cost	40%	20%	0%	5%	0.5%
Rare Trading, Low Cost	40%	20%	0%	1%	0.5%
Rare Trading, High Cost	40%	20%	0%	1%	1.0%
Nonlife Insurance Companies					
Benchmark	35%	35%	5.25%	0%	0.0%
Rare Trading, Low Cost	35%	35%	5.25%	5%	0.2%
Medium Trading, Low Cost	35%	35%	5.25%	10%	0.2%
Frequent Trading, Low Cost	35%	35%	5.25%	15%	0.2%
Medium Trading, High Cost	35%	35%	5.25%	10%	0.4%

Table 1: Different investor types face different tax codes (τ, τ_G, τ_E), propensities to trade λ, and costs of trading K.
Note that the tax rate on tax-exempt income τ_E is nonzero for some investors.

For all these reasons, different investors may face different tax codes (τ, τ_G, τ_E), propensities to trade λ, and costs of trading K. In order to study the sensitivity of the results to several assumptions, I derive the optimal trading strategy for nine different investor types. Table 1 lists the types. In the rest of the paper, however, for brevity's sake, I will focus only on the four cases that I deem to be most representative of actual price-setting investors in tax-exempt bonds: the benchmark case ($\lambda = 0$ and $K = 0$ for both individuals and nonlife insurance companies), and a case with moderate trading and transaction costs ($\lambda = 1\%$ and $K = 0.5\%$ for individuals, and $\lambda = 5\%$ and $K = 0.2\%$ for nonlife insurance companies).

2.6 The investor's problem, or pricing the bond

The bond matures at time T. At time $t < T$, a risk-neutral investor maximizes the present value of after-tax final wealth. The investor's utility is

$$V_t\left(c,\, T-t,\, B_t,\, r_t\right),$$

a function of the constant coupon rate c and three time-varying state variables: the short rate r_t, the time to maturity $T-t$, and the investor's tax basis B_t.

At every period, the investor can choose to do nothing ("hold"), or to realize the position's unrealized gain or loss by selling the bond and buying it back ("sell"). If the investor decides to hold, however, the bond will be sold (and the gain or loss realized) with a probability λ. This yields the following definition for the value function:

$$V_t\left(c,\, T-t,\, B_t,\, r_t\right) = \max_{i \in \left\{\text{sell, hold}\right\}} \frac{c + \mathbb{E}_t\left[V_{t+1}^i\right]}{1+r_t} - \tau_G\left(P_t - B_t\right)\mathbf{1}_{\left(i=\text{sell}\right)} \tag{1}$$

where P_t is the market price of the bond, c is the coupon, $\mathbb{E}_t\left[V_{t+1}^i\right]$ the next-period residual value, and $\mathbf{1}_{\left(i=\text{sell}\right)}$ is an indicator function that is 1 if the investor decides to sell and 0 otherwise. Moreover,

$$V_{t+1}^{\text{sell}} = V_{t+1}\left(c,\, T-t+1,\, B\left(P_t,\, c,\, T-t\right),\, r_{t+1}\right) \tag{2}$$

$$V_{t+1}^{\text{hold}} = \lambda V_{t+1}^{\text{sell}} + \left(1-\lambda\right) V_{t+1}\left(c,\, T-t+1,\, B\left(B_t,\, c,\, T-t\right),\, r_{t+1}\right) \tag{3}$$

and finally

$$B\left(x,\, c,\, T-t\right)$$

is the "constant yield method" function that gives the next-period tax basis given the current basis x, the coupon c, and the time to maturity $T-t$. Because of the need to solve numerically for the yield, this function cannot be expressed analytically, but it is altogether very simple. Note that in V_{t+1}^{sell} the first argument to $B\left(\cdot\right)$ is P_t, because selling the bond resets the basis equal to the current price.

So far we have taken the existence of a market price P_t for granted. However, V_t depends on P_t; before we derive the optimal policy (hold or sell) at the current period, we need to calculate the current period market price. The price is defined recursively as a special case of the value function:

$$P_t\left(c,\, T-t,\, r_t\right) = V_t\left(c,\, T-t,\, P_t,\, r_t\right). \tag{4}$$

Note that P_t appears on both sides! If the investor buys the bond now, the basis now will be equal to the price. Of course, the value of the position depends on the tax basis, so the price is a function of itself. As shown in Appendix C, this presents no practical problem even if no analytical solution exists, because $V(\cdot)$ is a contraction map with respect to its third argument. Any guess for P_t converges in a few iterations to the unique fixed point, and thus the price P_t can be recovered by just knowing the interest rate r_t, the time to maturity $T-t$, and the value function at the next period V_{t+1}. Once obtained P_t, it is easy to calculate V_t for all other values of the tax basis B_t using (1).

2.7 The issuer's problem

The issuer's problem is simple. $P_t\left(c,\, T-t,\, r_t\right)$ as defined in the previous subsection is the issue price of a bond maturing in $T-t$ years with coupon c when the current price is r_t. Define $\hat{P}_t$ as the cost to the issuer of all the future cash flows from the bond:

$$\hat{P}_t\left(c,\, T-t,\, r_t\right) = \mathbb{E}_t\left[\frac{c}{1+r_t} + \frac{c}{\left(1+r_t\right)\left(1+r_{t+1}\right)} + \cdots + \frac{1+c}{\left(1+r_t\right)\cdots\left(1+r_{T-1}\right)}\right]$$

The issuer wants to maximize the subsidy ratio W, i.e. the share of issued amount that is not compensation for future cash outflows:

$$\max_{c\geq 0} W \equiv \frac{P_t - \hat{P}_t}{P_t}$$

The difference now is that W also depends on the cost of trading K (which the issuer won't have to bear directly) and all three tax rates $(\tau,\, \tau_G,\, \tau_E)$.

The optimal coupon c^* is simply found by solving the investor's problem for set values of $c \in \{.1\%,\, .2\%,\, ...\, 20\%\}$ and picking the one that gives the highest W. This rough "grid search" method yields a less-precise optimum compared to a hill-climbing or annealing algorithm. However, given the nature of the problem, this level of precision is good enough. Moreover, intuition from a simple two-period model reported in Appendix D suggests that a finite optimum may not always exist. In these cases, "grid search" will simply yield $c^* = 20\%$, the maximum value available, while more sophisticated algorithms will fail. In practice, an "interior" optimal coupon (i.e. below 20%) exists very often.

3 Results: optimal trading strategy

In this section, I describe the optimal trading strategy for the investor. I also provide a simple approximate sell signal and show that it performs very well.

3.1 Description

The optimal trading strategy is visualized in Figures 5 (assuming no transaction costs) and 6 (assuming positive transaction costs). The purpose of this subsection is to explain the figures in detail. Subsection 3.2 shows some practical examples of trading decisions an investor may face, and introduces an approximate method that recovers the optimal strategy with little effort and good precision.

Each plot in the figures has time to maturity on the horizontal axis and tax basis on the vertical axis. Each plot assumes a different value of the current interest rate r_t (2%, 4.5%, 5.5% and 8% respectively), and all plots refer to a 5 percent coupon bond. Thus, each point on a plot represents a state vector $(t,\ r_t,\ B_t)$ and can be associated with an optimal trading strategy: "sell" (white) or "don't sell" (gray).

The strategy consists of up to three thresholds, separating "sell" areas from "don't sell" areas. ("Sell" here indicates a sell-and-buy-back transaction, in which the investor realizes unrealized gains or losses but the bond does not actually change hands.) The first threshold is the price (thick dark line). Because the bond's trading price is a function only of time to maturity $T - t$ and current short rate r_t, the price can be drawn on each plot as a function of time. Points above the line represent unrealized losses, because the investor's tax basis is higher than the price; viceversa, points below the line represent unrealized gains.

The second threshold is the "Time Value of Money" wall (TVM Wall, thin dashed line). Because of the specific tax rules, if maturity is far enough in the future, it is profitable to realize losses below par, as the benefit is immediate and the cost is low in present value because it is spread over a long period of time. However, if maturity is close enough, it is profitable to realize gains below par, precisely for the opposite reason - the cost is now, but the benefit will be realized quickly enough and it is larger in present value.

The third threshold is the "Not-Enough" wall (thin dotted line). Sometimes there are trans-

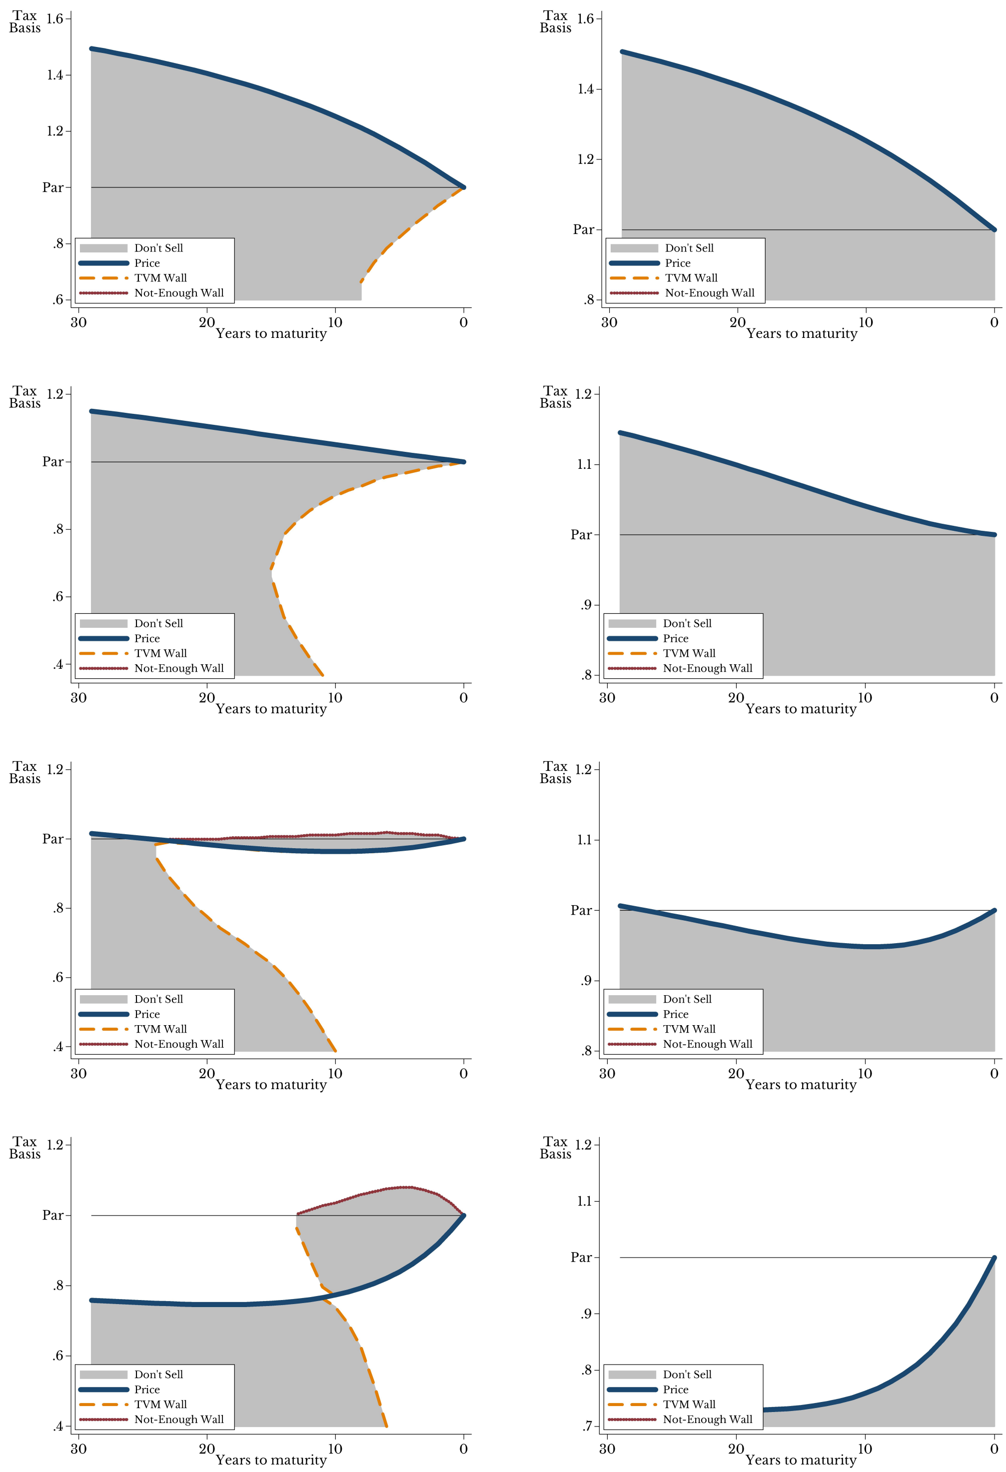

Figure 5: Optimal trading strategy with no transaction costs.
Each row corresponds to a different interest rate (r_t = 2%, 4.5%, 5.5%, 8%); left column is for individual investors, right for property and casualty insurance companies. (See explanation in text).

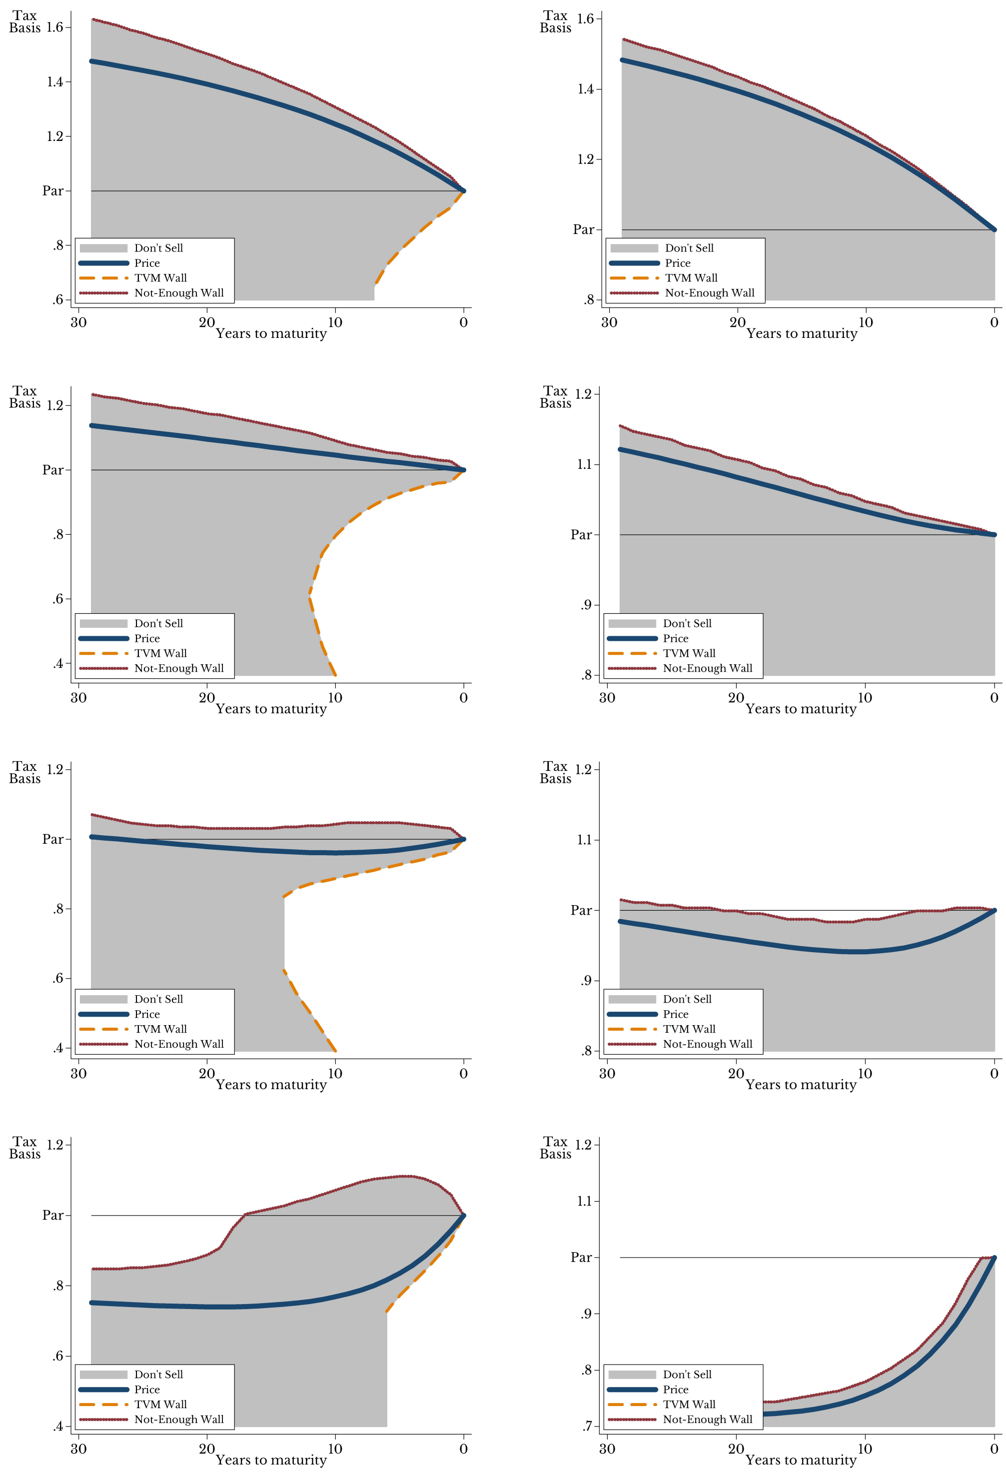

Figure 6: Optimal trading strategy with transaction costs.
Each row corresponds to a different interest rate (r_t = 2%, 4.5%, 5.5%, 8%); left column is for individual investors, right for property and casualty insurance companies. (See explanation in text).

action costs; sometimes, realizing losses above par is profitable, but below par it is costly; or realizing gains below par is profitable but gains above par are costly. “Not-Enough” means that the investor should not sell because gains or losses of the good type are not enough to make up for the transaction costs (if any) and for the cost of realizing the gains or losses of the bad type.

3.2 Practical examples

Properly read, Figures 5 and 6 are useful to gain economic intuition about the optimal strategy – its general form and the underlying tradeoffs. In this subsection I examine several practical examples with actual numbers and I relate them to specific plot points from Figures 5 and 6, to aid the understanding of the rather abstract explanation in the previous subsection.

Take for instance the bottom-left plots of Figures 5 and 6, showing the optimal strategy for an individual investor facing respectively no trading frictions ($K = 0$, $\lambda = 0$) and some trading frictions ($K = 0.50\%$, $\lambda = 5\%$), holding a 5% coupon bond when the current interest rate is 8%. By definition, a 5% bond is a discount bond when the interest rate is 8%: in both Figures, the price (indicated by the thick line) is below the horizontal axis (“Par”).

		$\lambda = 0$, $K = 0$ (Fig. 5)		$\lambda = 5\%$, $K = 0.50\%$ (Fig. 6)	
T	B	P	Sell?	P	Sell?
20	1	0.746	Realize	0.737	Realize
10	1	0.773	Don't	0.766	Don't
5	1	0.839	Don't	0.834	Don't
20	1.07	0.746	Realize	0.737	Realize
10	1.07	0.773	Realize	0.766	Realize
5	1.07	0.839	Don't	0.834	Don't
20	1.1	0.746	Realize	0.737	Realize
10	1.1	0.773	Realize	0.766	Realize
5	1.1	0.839	Realize	0.834	Don't

Table 2: Optimal trading strategy for an investor holding a 5% coupon bond of remaining maturity T with tax basis (book value) B.
Each row of the table (a T, B pair) represents one point on the bottom-left plot (individual investor, $r = 8\%$) of Figure 5 ($\lambda = 0$, $K = 0$: no transaction costs or liquidity shocks) or Figure 6 ($\lambda = 5\%$, $K = 0.50\%$: no transaction costs or liquidity shocks). For each situation the table reports price P and “sell signal”, taken directly from Figure 5. If point (T, B) is white (gray), the exact strategy recommends to realize (not to realize) the loss.

Table 2 analyzes several situations an investor may face. Each row of the table represents a bond of remaining maturity T with tax basis (book value) B. Each row in the table corresponds to a unique (T, B) pair, i.e. to a unique point on a plot (left panel for the bottom-left plot of Figure 5, without trading frictions; right panel for the corresponding plot of Figure 6, with trading frictions).

For each (T, B) pair, the table reports the strategy based on the dynamic programming solution: "Realize" gain or loss if the point on the plot is white, and "Don't" if gray.

The first block of three rows has $B = 1$ (tax book value equal to par value) and time to maturity $T \in \{5, 10, 20\}$. Both with and without transaction costs, point $(10, 1)$ falls within the gray "don't sell" area to the right of the "Time Value of Money Wall". On the other hand, point $(20, 1)$ falls in the white "sell" area to the left of the TVM Wall and realizing losses is profitable.

The second and third blocks assume a tax book value larger than par value: $B = 1.07$ and $B = 1.1$, respectively. The higher the tax book value, the more advantageous it is to realize losses, because a part of the loss ($B - 1$, i.e. 7 and 10 cents respectively) will give rise to a loss deduction now, and no tax liabilities in the future. When $B = 1.07$, for instance, the investor optimally realizes the loss when $T = 10$; when $B = 1.1$, the investor realizes the loss even when $T = 5$ in the case with no transaction costs (left).

The point $(5, 1.07)$ is interesting because it is below the "Not-Enough" threshold: in spite of the fact that a part of the loss (7 cents) can be realized at very advantageous conditions, the rest of the losses is too costly to realize, and overall the investor optimally holds on to the bond. Similarly, the point $(5, 1.1)$ is just above the "Not-Enough" threshold in the case with no transaction costs, and just below the threshold in the case with transaction costs.

3.3 Value of the tax option for the investor

Figure 7 plots the investor's gain from following the optimal trading strategy (realizing gains and losses optimally) compared to a buy-and-hold strategy. The gain is approximately 15 basis points per year for individual investors, and 25 for insurers. This gain is large for a risk-free fixed-income strategy. Stock trading strategies such as value or momentum promise annual "alpha" that is an order of magnitude larger; however, tax trading of tax-exempt bonds produces "alpha" that is based on ex ante theory, and therefore not subject to "model risk," and risk free, because it is generated through tax arbitrage. Quantitative stock strategies are actually very risky, even if the model is correct, and even though they may compensate very

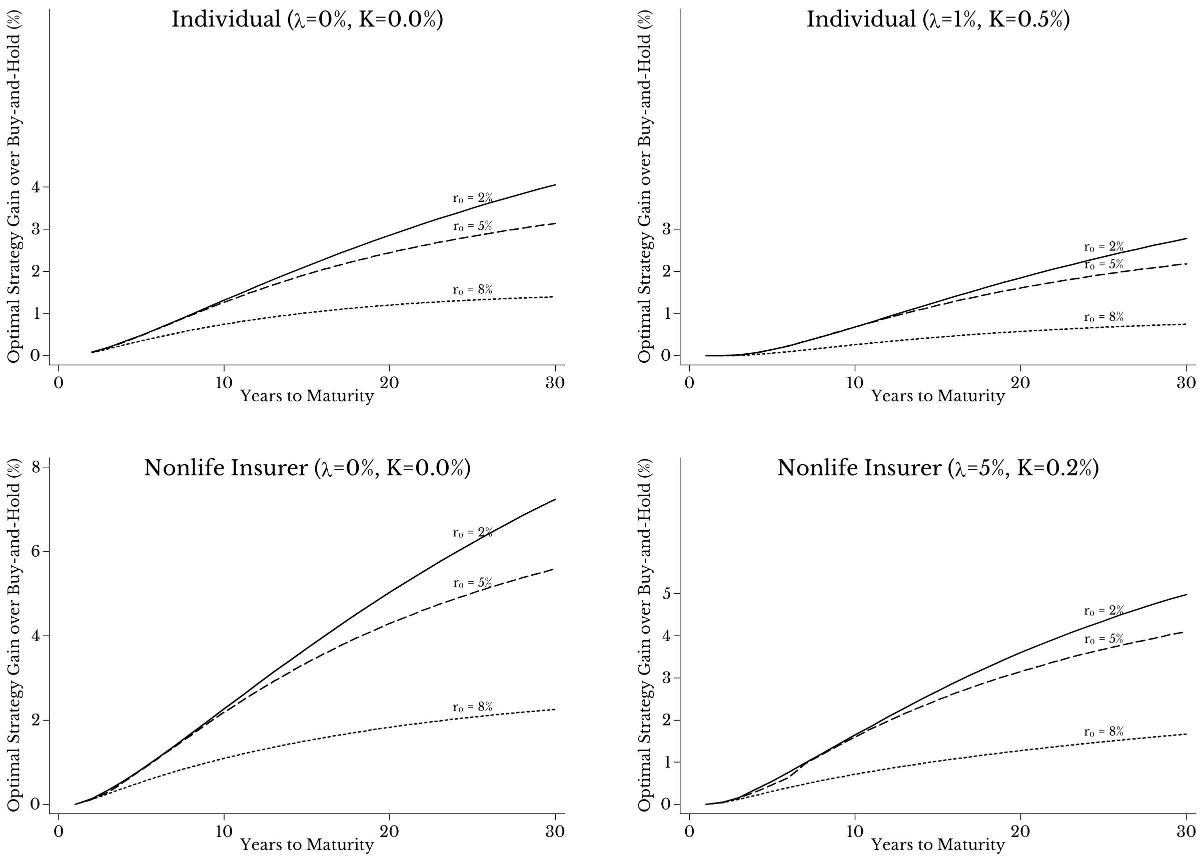

Figure 7: Tax option value as a fraction of issue price.
By realizing gains and losses optimally, investors can obtain a considerable improvement over a buy-and-hold strategy.

generously for this risk.[17] Finally, tax trading does not require short selling, or the purchase of a myriad securities. The strategy can be carried out by the owner of even a single bond.

4 Results: optimal issuance strategy

In this Section, I describe the optimal issuance strategy for the issuer and compare it with summary data on actual issuance.

[17] The optimal trading strategy does entail some change in the timing of cash flows, but an investor who is averse to interest rate risk can hedge it very cheaply using forward contracts. Moreover, tax rate risk, i.e. the possibility that the same individual faces different tax rates at different points in time, is a minor issue. Most of the option value comes from realizing losses, so that the benefits materialize immediately at a known tax rate, and the costs materialize in the future, if ever. For these future costs, I have assumed that the investor is in the top tax bracket, and therefore it is unlikely that the tax rate will be higher in the future, unless the government actually increases the tax rates. For the current benefits, for individuals, I have conservatively assumed that losses are deducted at the capital gain tax rate, while in actuality some losses can be deducted from ordinary income at a more favorable rate. In general, having time-varying tax rates creates more, not less, option value for taxpayers.

4.1 Qualitative features of optimally issued bonds

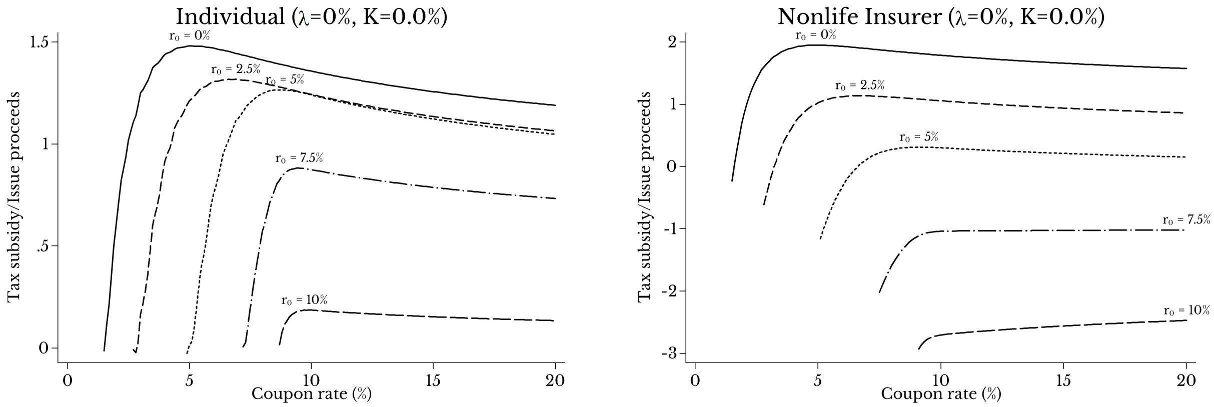

Figure 8: Value of the subsidy ratio W.
Value of the tax subsidy to issuers as a fraction of total amount issued, for a 10-year noncallable tax-exempt bond when the marginal issue buyer is an individual (left) or a nonlife insurance company (right). The leftmost point of each curve represents par issuance. The value increases rapidly at first, as more and more future potential losses become tax-deductible; then it tapers off (in most cases), as nearly all losses are already tax deductible, but further increasing the coupon decreases price volatility and hence option value. Note that for nonlife insurance companies tax-exempt interest is taxed at a positive rate of 5.25%, and hence the subsidy can be negative. This is not problematic; conditional on issuing, the optimal coupon rate maximizes the subsidy, regardless of sign.

Premium is optimal Figure 8 plots W (in percentage points) over the set of c values enumerated in Subsection 2.7 above (except for those values of c that are known to be suboptimal because the issue price is below par). The leftmost point of each curve represents par or near-par issuance. The marginal value of increasing the coupon is very high at first (a very steep hill), as more and more future potential losses become tax-deductible. At the optimum, most future losses have become tax deductible. From now on, in most cases, the marginal value of increasing the coupon is negative: increasing the coupon lowers the duration of the bond, decreasing price volatility, and with it, the value of the investor's tax option. (All this is in relative terms: both dollar price and dollar option value are increasing in the coupon, but after a while price increases faster than option value because of the effect described). It is worth noting that the optimal coupon rate is often below 10%, the maximum allowed interest rate; therefore, the existence of a unique optimum does not depend on the mechanical fact that the interest rate is bounded above for computational purposes.

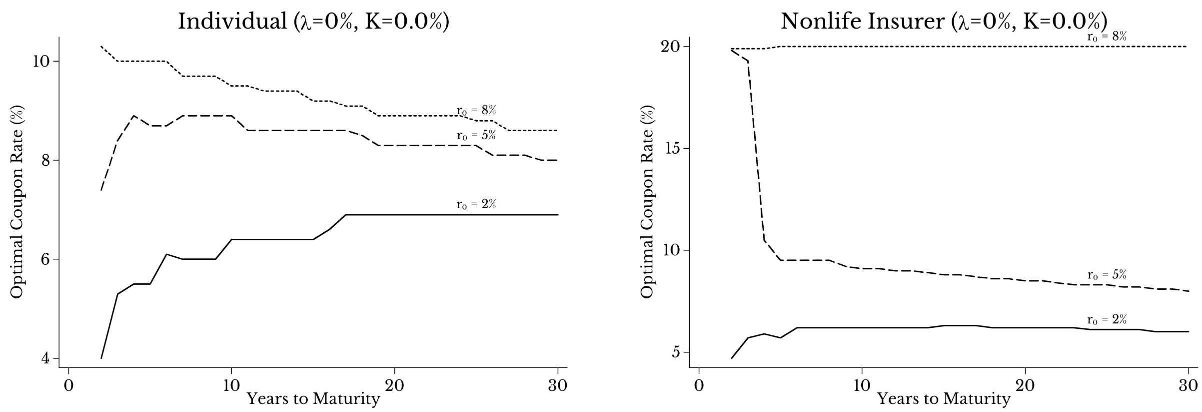

Figure 9: Optimal coupon as a function of bond maturity.
For individuals, the optimal coupon is a finite number in every scenario and for every maturity. For nonlife insurers, however, the optimal coupon is 20% or larger (possibly infinite) when the interest rate is high enough. As the bond maturity increases, the sensitivity of the optimal coupon to the current interest rate drops considerably, i.e. coupons are "sticky". The presence of "steps" in the graph is an artifact of the coarse grid search optimization described in Subsection 2.7.

Coupons are "sticky" This is an interesting case of the prediction that the optimal coupon is increasing in the expected change in yields. The interest rate process assumed for this Section's dynamic programming exercise exhibits a degree of mean reversion, like most widely used interest rate processes. Thus, when the interest rate is low, one expects it to rise, i.e. one expects the bond to have more losses than average, and a higher issue price (a higher coupon) is needed to make sure all these losses are deductible. The reverse happens when interest rates are high. The net result is that even when interest rates fall to a very low level (e.g., 0%), the optimal coupon stays high (e.g., 5%). Figure 9 shows how "stickiness" varies with maturity and the level of interest rates. For the high-interest rate scenario ($r_0 = 8\%$) for insurance companies (right), the optimal coupon is higher than 20%—and possibly infinite—and therefore no stickiness exists.

Term structure of issue prices Figure 10 displays the "term structure of issue prices." For both individual and corporate investors, the optimal issue price is an increasing and concave function of bond maturity. This is a consequence of the simple principle that in order to ensure that potential future losses are fully tax deductible, price should never fall below par. Other things equal, a longer-term bond has a more volatile price, and a higher issue price is needed to accommodate this volatility.

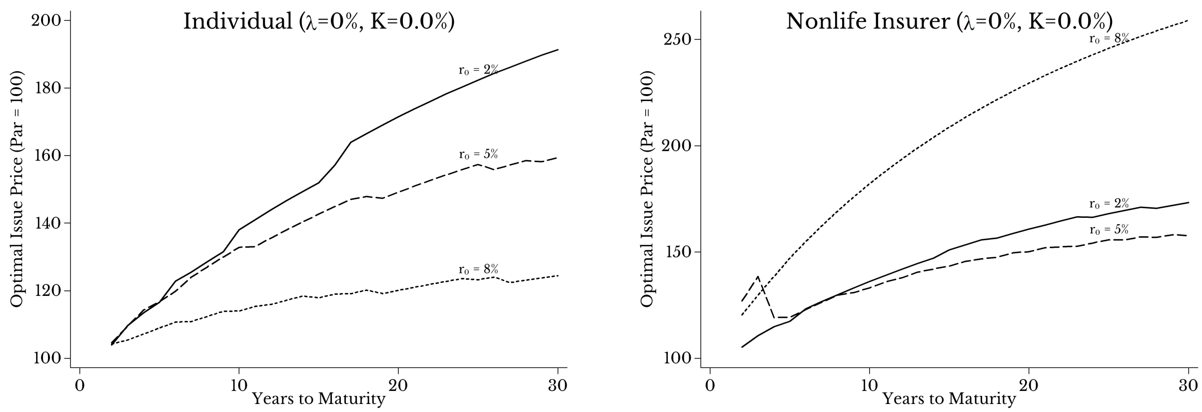

Figure 10: The "term structure of issue prices."
Optimal issue prices (expressed in percentage points of par) are an increasing and concave function of bond maturity. To ensure that potential future losses are fully tax deductible, price should never fall below par. Other things equal, a longer-term bond has a more volatile price, and hence requires more "protection" in the form of a higher issue price. The presence of "steps" in the graph is an artifact of the grid search optimization described in Subsection 2.7.

4.2 Value of optimal issuance

Figure 11 plots the issuer's gain from issuing a bond with the optimal coupon, compared to the benchmark case of issuing at par. For instance, if the ex-ante value of the tax option is 3 percent of the issued amount when the bond is issued optimally, and 1 percent when the bond is issued at par, then the gain is 2 percent.

From the Figure it is evident that once again, the benefit is large. The gain from issuing optimally is between 0.5 and 2 percent of the principal amount issued. To put this number in context, the cost of issuance itself is in the order of 0.25 to 1.5 percent of the amount issued: therefore, by setting the coupon optimally, a tax-exempt issuer can sometimes more than recoup the cost of issuing, thanks to the "tax option" value created at the expense of the U.S. Treasury.

4.3 Evidence: The features of actual tax-exempt bonds

In this Subsection I compare the qualitative features of optimally issued bonds obtained from the model and those of real-world bonds.

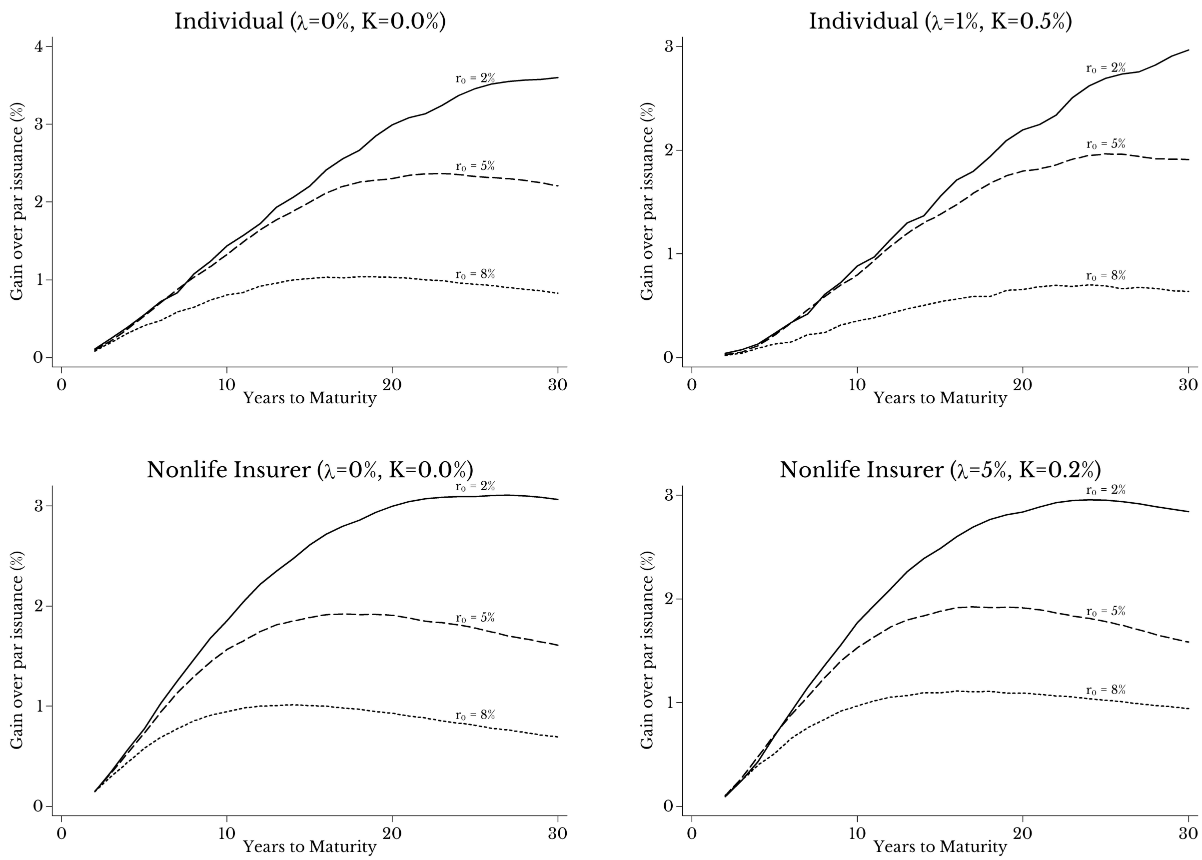

Figure 11: Tax option value as a fraction of issue price: optimal coupon improvement over par issuance.

4.4 Data

The data are a combination of information from Bloomberg and the National Association of Insurance Commissioners (NAIC). NAIC filings provide a quarter-by-quarter history of all bond positions of all insurers. I selected the top 2,100 property and casualty insurers, corresponding to about 99.5% of all bonds by dollar value of positions, and I took all unique CUSIP numbers, about 200,000 unique bonds. The data was available from 2004q4 to 2012q3 at the time of acquisition. For each of these bonds, Bloomberg provides a near-official issue price, whether the bond has an issuer call option, and when it can be exercised; and the taxable status of the security (because multiple classes of municipal bonds are taxable).

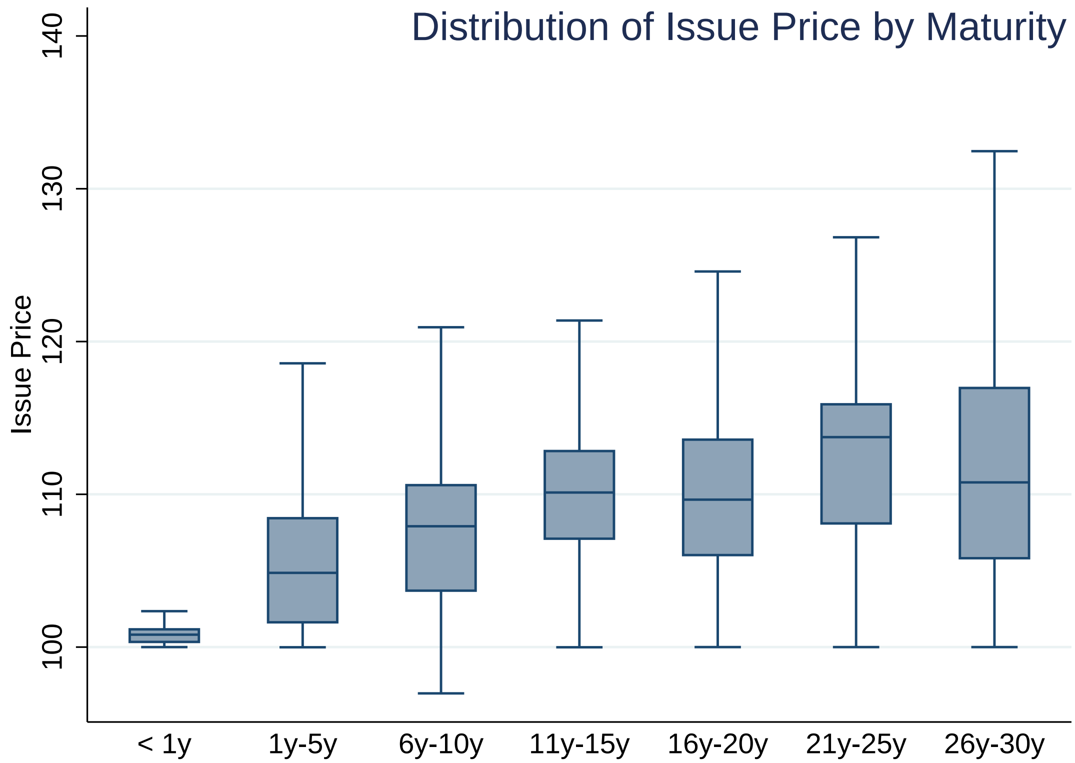

Figure 12: Distribution of issue premium by maturity.
The box contains 50% of the empirical distribution; only a few outlying datapoints are outside the area enclosed by the whiskers. The stylized model of Section D predicts that issue price of noncallable tax exempt bonds is increasing in expected price volatility. This figure confirms that prediction. Bond maturity is a natural driver of ex-ante expected volatility for noncallable bonds. This pattern is very robust, and it emerges within every single vintage of issuance.

	Corporate	Government	Muni (Exempt)	Muni (Taxable)
Maturity	-0.04***	-0.06***	0.43***	-0.02***
	(-4.95)	(-4.76)	(70.54)	(-5.51)
Maturity × Callable	-0.07***	-0.43***	-0.15***	-0.01***
	(-14.12)	(-34.58)	(-44.59)	(-4.11)
Callable	0.09	2.68***	2.76***	-0.09
	(0.87)	(15.40)	(38.66)	(-1.71)
Constant	100.05***	100.25***	103.99***	100.49***
	(1313.80)	(817.04)	(1942.32)	(3297.32)
N. Obs.	15,964	11,354	98,143	8,280
R squared	0.016	0.098	0.142	0.009

* p<0.05, ** p<0.01, *** p<0.001. t-statistics in parenthesis.

Table 3: Issue prices of tax-exempt bonds, unlike other bonds, are increasing in time to maturity.

4.5 Optimality of premium and the term structure of issue prices

A striking confirmation of the model's predictions is found in Figures 1 and 12. Municipal bonds are indeed issued at large premiums to par, in a sample including bonds issued throughout 20 years, and the term structure of issue prices is increasing and concave in the maturity of the bonds for noncallable bonds. More formal evidence is given in Table 3.

Actual issue prices (and coupons) are lower than the optimal ones predicted by the dynamic programming exercise. The model in this paper takes into account taxes only; perhaps, there are other factors that affect the optimal coupon, and further research is warranted. This is good news, however, because it means that taxes alone would be enough to explain even larger issue premiums; the opposite pattern (actual coupons higher than optimal) would have been much harder to justify.

An important caveat regards the inherent sample selection problem. Insurance companies may be more tax-conscious than the average individual, and thus this sample of bonds may exaggerate the strength of the findings. In any case, this is either evidence that issuers issue intelligently designed securities, or that insurance companies choose carefully what securities to buy.

5 Conclusion

In this paper I show that after-tax pricing of tax-exempt bonds has security design implications for tax-exempt issuers. In a stylized model of optimal issuance, I derive that (i) tax-exempt bonds should be issued at a premium, and (ii) issue price should be increasing in the expected volatility of the bonds and in the expected change in yields—all to maximize the future gain and loss harvesting opportunities for the investor. I refine these qualitative predictions using dynamic programming to price tax-exempt bonds: with some degree of mean reversion in yields, I generate the additional prediction that coupons are "sticky", i.e. they stay high when yields fall. I show empirical evidence consistent with tax options being priced, and with issuers issuing optimally designed securities. This is striking evidence that at least some institutional bond buyers are very sophisticated, and that market prices transmit information efficiently even in an opaque and illiquid market such as that for municipal bonds.

References

Amiram, D., Kalay, A., Kalay, A., Ozel, N. B., 2014. Capital market frictions and bond coupon choice, columbia Business School Research Paper No. 13-13, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2230078.

Anderson, R. W., Sundaresan, S., 1996. Design and valuation of debt contracts. Review of Financial Studies 9(1), 37–68.

Ang, A., Bhansali, V., Xing, Y., 2010. Taxes on tax-exempt bonds. Journal of Finance 65(2), 565–601.

Bhamra, H. S., Kuehn, L.-A., Strebulaev, I. A., 2010. The aggregate dynamics of capital structure and macroeconomic risk. Review of Financial Studies 23(12), 4187–4241.

Brown, D. T., 2011. Callable bonds: What is special about munis?, available at SSRN: http://ssrn.com/abstract=1935647.

Chalmers, J. M. R., 1998. Default risk cannot explain the muni puzzle: Evidence from municipal bonds that are secured by u.s. treasury obligations. Review of Financial Studies 11, 281–308.

Constantinides, G. M., Ingersoll, Jr., J. E., 1984. Optimal bond trading with personal taxes. Journal of Financial Economics 13, 299–335.

DeMarzo, P. M., Fishman, M. J., 2007. Optimal long-term financial contracting. Review of Financial Studies 20(6), 2079–2128.

Elton, E. J., Green, T. C., 1998. Tax and liquidity effects in pricing government bonds. Journal of Finance 53(5), 1533–1562.

Green, R. C., 1993. A simple model of the taxable and tax-exempt yield curves. Review of Financial Studies 6, 233–264.

Green, R. C., Ødegaard, B. A., 1997. Are there tax effects in the relative pricing of U.S. government bonds? Journal of Finance 52(2), 609–633.

Jordan, B. D., 2012. On the relative yields of taxable and municipal bonds: A theory of the tax structure of interest rates. The Journal of Fixed Income Summer, 79–87.

Jordan, B. D., Jordan, S. D., 1991. Tax options and the pricing of Treasury bond triplets—theory and evidence. Journal of Financial Economics 30, 135–164.

Kalotay, A., 2013. Optimum tax management of municipal bonds, working Paper (on file with the author).

Kalotay, A., Howard, C. D., 2014. The tax option in municipal bonds. Journal of Portfolio Management forthcoming.

Litzenberger, R. H., Rolfo, J., 1984. Arbitrage pricing, transaction costs and taxation of capital gains. A study of government bonds with the same maturity date. Journal of Financial Economics 13, 337–351.

Longstaff, F. A., 2011. Municipal debt and marginal tax rates: Is there a tax premium in asset prices? The Journal of Finance 66(3), 721–751.

Trzcinka, C., 1982. The pricing of tax-exempt bonds and the miller hypothesis. The Journal of Finance 35(4), 907–923.

A The taxation of bonds

This section is meant to familiarize the reader with the relevant U.S. federal income tax rules (throughout the paper, I ignore state and local taxes for simplicity). Readers may skip this section and come back for reference later, should any details or modeling choices turn out to be unclear.

While this paper is exclusively concerned with tax-exempt bonds, here I discuss the taxation of taxable and tax-exempt bonds together; this entails little extra cost and helps to more clearly highlight the differences that drive all the results in this paper.

Figure 13 is meant as a guide and summary to this section. In each of the six charts, a point represents a transaction. The coloring of each area indicates what tax regime applies to the transaction, i.e. to the points that fall into that area. Thicker polka dots represent a greater tax incentive to sell. As the Figure shows, the taxation of bonds creates *five* distinct tax regimes! Moreover, as explained below, more than one regime could apply to a given sale.

The horizontal axis represents the time of the transaction (0 = time of issue), and the vertical axis represents the sale price (100 = face value). Bonds can be taxable (left) or tax-exempt

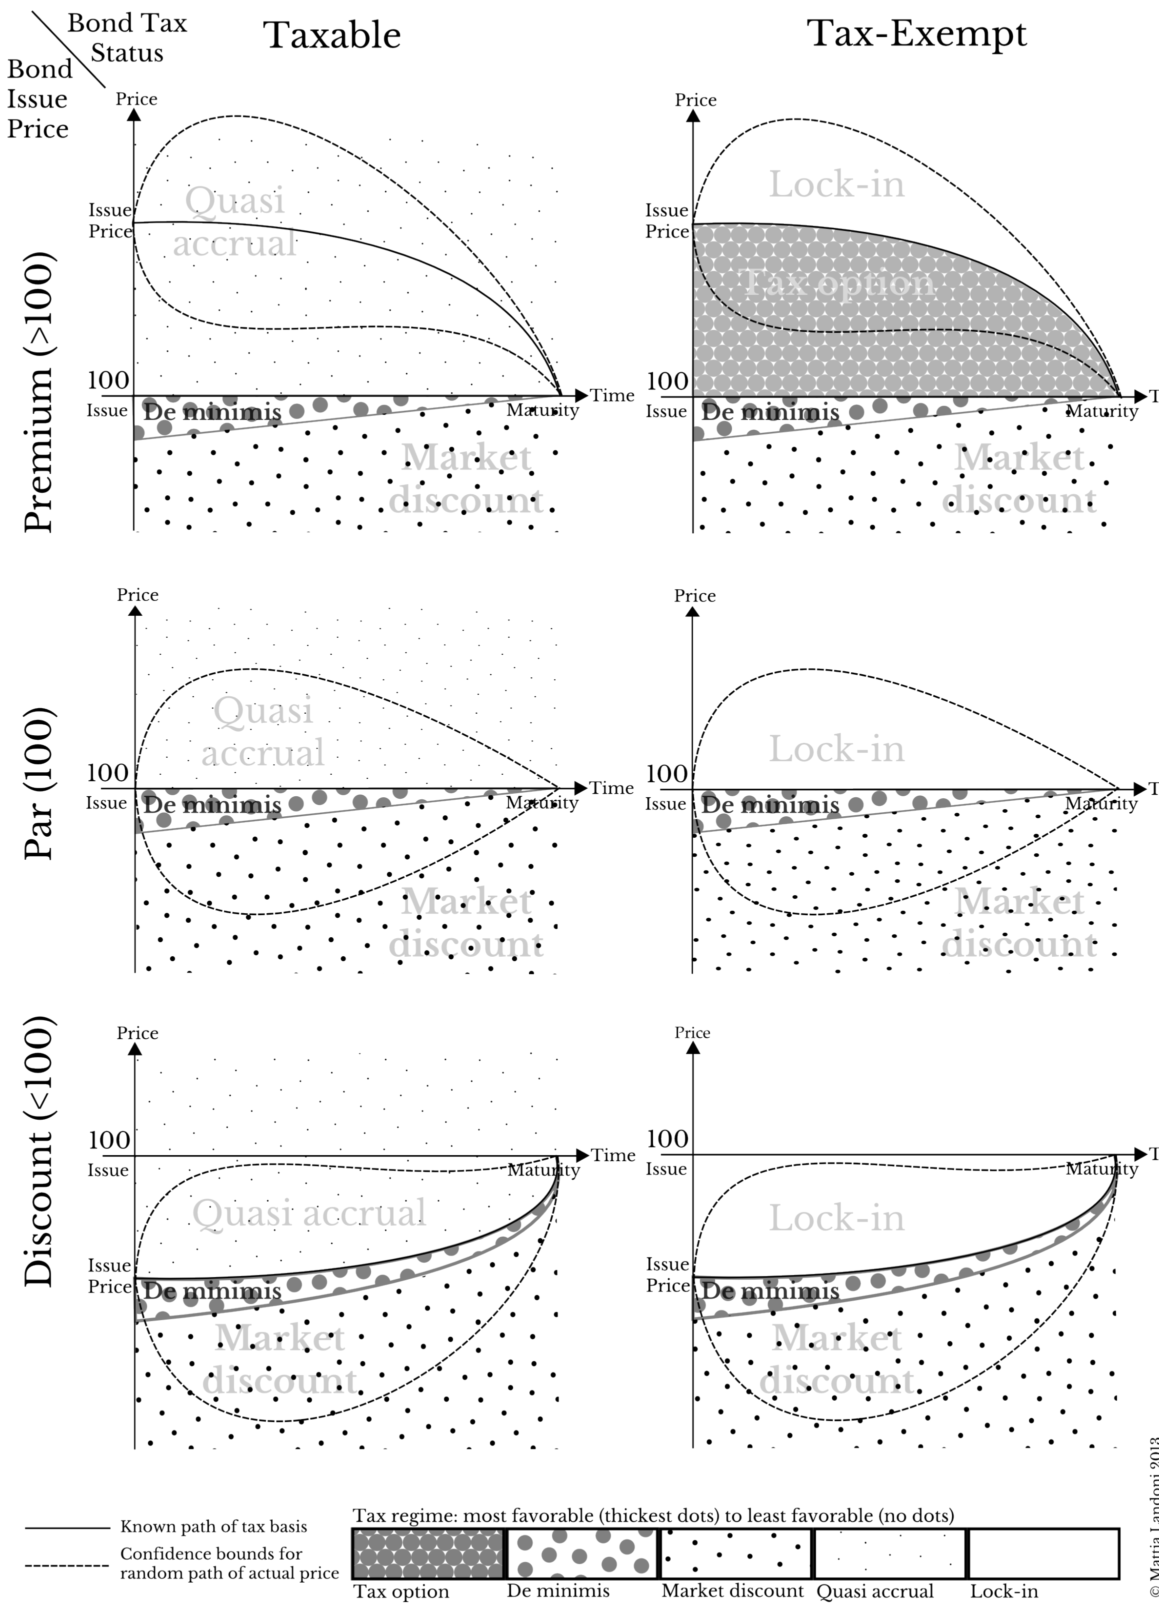

Figure 13: Two types of bonds, five tax regimes.
(See text for an explanation)

(right), and they are issued at a price equal to face value (par), below (discount) or above (premium).

The solid black line is the predetermined time path of the tax basis for an investor who bought the bond at the issue price, and receives 100 (the bond's face value) when the bond matures.[18] Accordingly, the tax basis starts at the issue price at time 0 (a point on the vertical axis), and gradually converges to 100 (a point on the horizontal axis) by the time of maturity. The dashed lines drawing an "avocado" around the tax basis are confidence bounds for the random path of the *actual* market price; unlike the basis, the market price is not known in advance.

A.1 Simplifying assumptions

For simplicity, Figure 13 and the following explanation assume that the transaction in question is the last one—there will not be any other transaction until maturity. As will be shown later, this has only a small effect on the overall tax incentive to sell.

In addition, to make the exposition less dense, the explanation in this section focuses on the tax consequences of a "sell-and-buy-back" transaction, i.e. a transaction in which the owner of the bond sells the bond and buys it back immediately at the same prevailing price. Focusing on a sell-and-buy-back, as opposed to a "traditional" sale, is without loss of generality. To understand why, consider the following thought experiment. Investor S holds a bond; her tax basis is lower than the current prevailing price, so by selling the bond S will realize a taxable gain. A buyer, B, comes along and makes an offer:

- S may sell the bond to B directly (a "traditional" sale). A gain or loss will be realized and the bond will change hands.
- Or, S can execute a sell-and-buy-back (all gain or loss will be realized, and the new basis will be equal to the market price), then sell the bond to B (the bond will change hands with no tax consequences, because no gain or loss is realized).[19]

Clearly then, all tax consequences that apply to a "traditional" sale are independent of the existence of a buyer distinct from the seller. Because of this "separability" of the tax and

[18] The tax basis is the book value of the asset for tax purposes: upon a sale, the seller realizes a capital gain or loss calculated as the difference between the sale price and the tax basis.

[19] The IRS disallows "wash sales", i.e. sell-and-buy-back transactions in which a emphloss is realized. Like in many other papers in the tax trading literature, in this paper I assume that this rule is not binding. An investor could, for instance, sell a bond and buy another that is very similar but not so similar to trigger the wash sale rule.

ownership consequences of a sale, if a sell-and-buy-back is advantageous for the investor from a tax standpoint, so is any sale; and viceversa.

Finally, throughout the paper I assume that gains and losses are taxed at the same rate τ_G, so realizing a one-dollar gain causes the investor to face a tax bill of τ_G, while realizing one dollar of loss the investor receives a rebate of $-\tau_G$. In reality, because of asymmetries in the treatment of gains and losses, one dollar of losses might be worth less than one dollar of gains in absolute value. This simplification is rather standard in the literature on taxes and trading.

A.2 Two types of bonds, five tax regimes

Tax law as it applies to income from bonds gives rise to five distinct tax regimes, which I call "lock-in", "quasi accrual", "de minimis", "market discount" and "tax option". Under all regimes, the tax bill assessed at the time of a sell-and-buy-back is identical. The investor pays capital gains taxes (if realizing a gain) or receives a loss deduction (if realizing a loss).

The difference between the regimes is entirely determined by what happens *after* the sell-and-buy-back. Tax law is such that every dollar of gain (loss) realized today is balanced by a one-dollar reduction (increase) in the taxable part of income flowing from the asset in the future. Each regime corresponds to a unique combination of (i) timing and (ii) applicable tax rates.

It is useful to classify the regimes depending on their reliance on an "economic" or a "cash" criterion.[20] For instance, take a bond issued at a price of 90, paying no coupon, and paying off 100 in two years. According to a cash criterion, interest income of 10 is realized in the second year, when the cash is received. According to an economic criterion, the bond yields a constant return of 5.41% per year, and thus the investor will record an income of 4.74 (5.41% of 90) in the first year, and 5.26 (5.41% of 94.74) in the second year. The assumption of a constant return is meant to more closely reflect the underlying economics of the situation.

Quasi accrual (economic) This regime applies to taxable bonds only—to all gains, and to any losses above face value (as in the top-left panel of Figure 13). Suppose an investor holds a bond with a tax book value of 120; the market price is 110. By executing a sell-and-buy-back, the investor posts an immediate taxable loss of 10 (120-110), and future taxable income from the bond is increased by 10. This extra income is spread over the remaining life of the bond as if it were extra yield. Under this regime it can be advantageous to

[20] The cash criterion is often called "realization", while the economic criterion is often called "accrual".

realize gains *or* losses, depending on the specific rates faced by the investor, and the time remaining to maturity; but the advantage or disadvantage is usually small.

De minimis **(cash)** This regime applies to all bonds whose current price is just below face value (for bonds issued at premium or par) or just below the original tax book value (for bonds issued at discount). Suppose the bond price is now 98.5, while the tax book value 120 as before. The investor posts a taxable loss of 21.5 (120-98.5). Of this, the first 20 (120-100) falls into the quasi accrual area so the investor will have 20 of extra ordinary income spread over the remaining life of the bond. The remaining 1.5 (100-98.5) falls into the *de minimis* area, so the investor will realize 1.5 of *capital gain*, but only when she disposes of the bond, or the bond matures.[21] Realizing *de minimis* losses with a sell-and-buy-back creates a taxable loss today in exchange for a taxable gain in the future. Because losses and gains are assumed to be taxed at the same rate, realizing *de minimis* losses is always advantageous.[22]

Market discount (cash) This regime applies to all bonds whose current price is anywhere below the *de minimis* threshold. Suppose the bond price is now 90, and the tax book value 120 as always. The investor posts a taxable loss of 30 (120-90). As before, the first 20 (120-100) falls into the quasi accrual area so the investor will have 20 of extra ordinary income spread over the remaining life of the bond. The remaining 10 (100-90) falls into the market discount area: when the bond matures or is disposed of, the investor will realize 10 of *ordinary income*. Thus, a sell-and-buy-back creates a taxable loss today in exchange for ordinary income in the future. Compared to the *de minimis* regime, income is taxed at the same time but at a higher tax rate, and thus the market discount regime is *less* advantageous; compared to the quasi accrual regime, income is taxed at the same tax rate but later, and thus the market discount regime is *more* advantageous. In absolute terms, realizing market discount losses may or may not be advantageous.

The next two regimes apply only to tax-exempt bonds. Together, they are the counterpart of the "quasi accrual" regime of taxable bonds.

Lock-in (economic) This regime applies to all gains on tax-exempt bonds. Suppose an investor holds a bond with a tax book value of 120; the market price is 130. Upon a sell-

[21] The *de minimis* treatment applies to bonds whose discount to face value is no greater than 0.25 times the number of complete years between the bond's acquisition date and its maturity date. For instance, a price of 98.5 qualifies for *de minimis* treatment only if the bond has at least six whole years left to maturity: $(100 - 98.5)/.25 = 6$.

[22] Here and throughout I assume no transaction costs for simplicity.

and-buy-back, the investor realized 10 of capital gains. As in the quasi accrual regime, future interest income is reduced by 10. However, interest income is already tax-exempt! The economic benefit from reducing future tax-exempt interest income is either null or very small, depending on the applicable rules, while the present cost of paying the capital gains tax is large and immediate. This regime makes selling very costly, and creates a so-called "lock-in effect".

Tax option (economic) This regime applies to tax-exempt bonds, and only to losses above face value: therefore, it appears only in the top-right panel of Figure 13. It is the opposite of "lock-in": realizing losses creates an immediate tax benefit, and increases the amount of future tax-exempt interest income, which does not entail a significant cost. As in the market discount example, suppose the bond price is 110 and the tax book value 120. The investor posts a tax loss of 10 (120-110). Future *tax-exempt income* increases by 10, spread over the remaining life of the bond as extra tax-exempt yield, and no additional tax is paid by anyone as a consequence of the transaction.[23]

B Differences between the tax code used in the dynamic programming exercise and the actual tax code

The tax code used in the dynamic programming exercise is relatively similar to the actual U.S. federal income tax code. For an investor who buys the bond for a price P_t, there are three scenarios:

1. $P_t = 1$: the bond is bought at par. In this case, the tax basis is 1 as long as the bond is held. All coupon cash flows constitute taxable income and are taxed at a rate τ_E. This rate is zero for individual investors, but it need not be zero for all investors.

2. $P_t > 1$: the bond is bought at a premium. In this case, the tax basis is equal to P_t when the bond is bought. The excess over face value $P_t - 1$ is called "premium" and it is amortized over the remaining life of the bond using the constant yield method.[24] The net income

[23]Generally, in the finance literature, every opportunity for profitable tax trading is called a "tax option". In the case of tax-exempt bonds, however, the value of realizing losses above par dwarfs all other opportunities for profitable tax arbitrage. Thus, for simplicity, throughout the paper "tax option" refers only to this option to realize losses above par.

[24]The constant yield method is described in detail by IRS Publication 550. In short, the taxpayer is to calculate the yield to maturity y at the time of purchase. Every year, the taxpayer multiplies y times the bond's book value (the tax basis). The difference between this amount and the coupon is the (potentially negative) amortization amount. The amortization amount should be added both to the tax basis and to income, but which type of income depends on the rules explained in the text.

(coupon minus amortization amount) is taxed at a rate τ_E.

3. $P_t < 1$: the bond is bought at a discount. In this case, the tax basis is equal to P_t when the bond is bought. The discount to face value $1 - P_t$ is amortized over the remaining life of the bond using the constant yield method. The amortization amount is added to the basis; the same amount is also recorded as market discount interest income and taxed at a rate τ. The entire coupon income is taxed at a rate τ_E.

The three differences concern the timing of market discount interest, accounting for original issue discount, and so-called "de minimis" discount.

Market discount Under the actual tax rules, market discount interest income is not taxed period by period as it accrues, but it is taxed when the bond matures or when the investor disposes of it. This adds one dimension to the dynamic programming problem, because one needs to keep track of *both* the tax basis *and* the purchase price. On the other hand, by assuming that market discount income is taxed as it accrues, it is sufficient to keep track of the basis. This can slightly change both the optimal trading strategy and the value of the bond. However, the intuition from the simple model in Appendix D suggests that rarely will an optimally issued bond trade at a market discount.

Original issue discount Under the actual tax rules, the threshold for market discount income is not always par (1). If the bond is issued with original issue discount (OID), i.e., the issue price is less than face value ($P_0 < 1$), the OID rules provide a way to calculate the boundary at every period. This is ignored, once again, to reduce the dimensionality of the dynamic programming problem; because otherwise one needs to keep track of the issue price. This assumption is especially harmless because, as always, issuing at a premium is optimal, and ignoring the OID rules affects only bonds issued at a discount, which are known ex-ante to be suboptimal.

De minimis discount Under the actual tax rules, when market discount is very small ("de minimis"), it is taxed at the capital gains tax rate. While this would not require any extra state variables, it makes things more complicated and its quantitative effect is, effectively, "de minimis". Therefore, this aspect can be safely ignored for the purposes of the present paper.

C Asset pricing with taxes using dynamic programming

Description of the problem

In a problem with taxes like the one encountered in this paper, the following situation often arises: an investor needs to decide how much to pay for an asset, but the price paid today affects the tax basis tomorrow, and therefore it affects future after-tax cash flows and therefore it affects the price itself. This Appendix shows that in essentially all cases, it is possible to find the price numerically using value function iteration.

Some iterative method appears to be used in practice by the few people who actually had to solve this problem, both academics such as Constantinides and Ingersoll (1984), and practitioners such as Kalotay Analytics (for after-tax pricing of municipal bonds) and PORTAX (for after-tax mean-variance optimization). While none of these people can be blamed for its absence, I am not aware of any published proof that iterative techniques should work under very general conditions.

Simple example

To understand what this means, suppose there is a bond trading today at a price P. Tomorrow, the bond pays back its face value plus a coupon: $1 + C$. Unfortunately, the interest income is taxable at a rate τ. Define interest income I as

$$I \equiv 1 + C - P$$

and thus the after-tax cash flow tomorrow is

$$1 + C - \tau\,(1 + C - P) = (1 + C)\,(1 - \tau) + \tau P.$$

In this instance, P is called the "tax basis" and it is directly evident that a higher tax basis today means more after-tax cash flow tomorrow.

However, P is also the price the investor pays today. If the applicable discount rate is r, the price of the bond is

$$P = \frac{1}{1 + r}\left[(1 + C)\,(1 - \tau) + \tau P\right]$$

In this case, it is easy to solve:

$$P=\frac{1-\tau}{1+r-\tau}\left(1+C\right)$$

However, often it is not practical to solve analytically for P (e.g., in the presence of uncertainty in discount rates, or in the timing or magnitude of cash flows). The next section defines and solves the problem in its most general form.

General result

Suppose that at time t the value of holding an asset position is

$$V\left(t,\ B_t,\ s_t\right)$$

where B_t is the tax basis and s_t is a random state vector that does not depend on P_t. Suppose also that the tax basis is equal to the purchase price at the time of purchase, and then it evolves deterministically according to some rules. Thus, the price an investor is willing to pay for an asset is equal to

$$P_t=V\left(t,\ P_t,\ s_t\right) \tag{5}$$

The question is whether it is safe to start with a guess for P_t, plug it into the right-hand side, and iterate until convergence. Here I try to show that this procedure is guaranteed to work, because $V\left(\cdot\right)$ is a contraction mapping with respect to its second argument, B_t. In order to show this, it is necessary to be a little more explicit on the form of $V\left(\cdot\right)$:

$$V\left(t,\ B_t,\ s_t\right)\equiv\frac{1}{1+r_t}\left(\mathbb{E}^Q\left[V\left(t+1,\ B_{t+1},\ s_{t+1}\right)+CF_{t+1}\left|s_t\right.\right]\right) \tag{6}$$

Where $\mathbb{E}^Q$ is the risk-neutral expectation operator. Then, make the following technical assumptions:

1. CF_{t+1} is the next-period pretax cash flow promised by the asset, and it does not depend on how much an investor pays for the asset at the current period.

2. $r_t\geq 0$: the one-period risk-free rate is nonnegative at all periods.

3. The top marginal tax rate τ is less than 100%.

4. The tax basis converges monotonically to a known finite number.

Then, it is easy to show

Proposition $\frac{\partial V(t,\,P_t,\,s_t)}{\partial P_t} < 1$

Proof The derivative of (6) with respect to B_t is

$$\frac{\partial V\left(t,\ B_t,\ s_t\right)}{\partial B_t} = \frac{\partial}{\partial B_t}\frac{1}{1+r_t}\left(\mathbb{E}^Q\left[V\left(t+1,\ B_{t+1},\ s_{t+1}\right)|s_t\right] + CF_{t+1}\right) =$$

$$= \underbrace{\frac{1}{1+r_t}}_{\leq 1} \quad \mathbb{E}^Q\left[\quad \underbrace{V_B\left(t+1,\ B_{t+1},\ s_{t+1}\right)}_{\leq \tau < 1} \quad \cdot \quad \underbrace{\frac{dB_{t+1}}{dB_t}}_{\leq 1} \quad |s_t\right] \quad < 1 \tag{7}$$

Assumption 1 guarantees that CF_{t+1} drops out when we take derivatives. Assumption 2, 3 and 4 guarantee the three bracketed inequalities of (7). Assumption 2 is self-explanatory. Assumption 3 guarantees the second inequality because one dollar of tax basis by definition reduces income from the asset by one dollar at some point in the future. If *all* the benefit from a one-dollar higher basis is realized *now* and at the *top marginal tax rate* τ, the benefit is τ. Assumption 3 guarantees that $\tau < 1$. Assumption 4 guarantees the third inequality because increasing the tax basis by one dollar today can increase the tax basis tomorrow by at most one dollar. Finally, since (7) holds for arbitary B_t, setting $B_t = P_t$ we obtain

$$\frac{\partial V\left(t,\ P_t,\ s_t\right)}{\partial P_t} < 1.$$

Proposition C, together with the fact that $\frac{\partial V(t,\,P_t,\,s_t)}{\partial P_t} > 0$, directly implies that $V\left(\cdot\right)$ is a contraction mapping, and therefore the unique fixed point can be found by iteration starting from any guess for P_t. In practice, the convergence is very fast.

There is a very good economic intuition between this result. Assumptions 1-4 really boil down to a simple and intuitive condition: that paying one dollar more for an asset today generates less than one dollar in present value of tax benefits. If this condition were violated in a real-world tax code, the government would go bankrupt very quickly and there would be no asset market equilibrium, as paying arbitrarily high prices for assets would make people richer (as long as the government is not bankrupt, that is). So, even if one can find examples where each one of the assumptions fails to hold,[25] the result is very general and robust.

[25] This is left as exercise for the creative reader.

A final thought on Blackwell sufficient conditions

One may try to construe this tax problem as a special case of policy function iteration; the Blackwell sufficient conditions are met (there is discounting, and the value function is strictly increasing in basis, hence in the price paid today). The hard-to-accept part is that the "policy" function would be a joint vector (P_t, γ_t) where γ_t is 0 when the asset is held, and 1 when a sell-and-buy-back is executed; and P_t is the price that the investor is willing to pay to buy an additional unit of the asset. But P_t is not "optimal policy" in any sense—optimal would be to pay exactly zero, if the investor could choose; and of course, the investor is assumed to be competitive, hence unable to choose.

D Optimal security design for tax-exempt bonds

This subsection has a little model that may be useful to some readers to build intuition about what drives the results in this paper.

An issuer of tax-exempt bonds (a state, local government or nonprofit institution) designs the bond in a way that maximizes the investor's future ability to generate valuable tax losses for federal tax purposes. A competitive investor pays the issuer for this "service", and in equilibrium this payment translates to an implicit fiscal subsidy from the national government to the subnational issuer.

D.1 Definitions and assumptions

There are two time periods ($t \in \{0, 1, 2\}$) and two agents (an issuer and an investor). The agents are risk-neutral and competitive, and there are no transaction costs. At time 0, the issuer issues a two-period bond with face value 1. In this simplified setup, I assume that the only security design choice available to the issuer is the coupon, c.[26]

Time value of money Both agents discount future payments at the same rate. One dollar received at time $t+1$ is worth δ_t at time t. I assume $\delta_t < 1$, i.e. the discount rate is positive. The

[26]Note that the *yield* is determined by market forces, but the *coupon* is the issuer's choice; if the coupon is higher than the yield, the bond will be issued at a premium to face value, and vice versa.

discount factor is random: while δ_0 is known at time 0, δ_1 becomes known only at time 1:

$$\delta_1 = \begin{cases} \bar{\delta} + \sigma & \text{with probability } p \\ \bar{\delta} - \sigma & \text{with probability } 1 - p \end{cases}$$

where σ can be interpreted as the volatility of the discount factor. Thus, the bond price at time 1 will be

$$P_1 = \delta_1 (1 + c) = \begin{cases} P_{\text{high}} \equiv (1 + c) \left(\bar{\delta} + \sigma\right) & \text{with probability } p \\ P_{\text{low}} \equiv (1 + c) \left(\bar{\delta} - \sigma\right) & \text{with probability } 1 - p \end{cases} \tag{8}$$

Taxes The issuer is tax exempt, and therefore the investor will not pay any tax on the coupon income. However, the investor does pay capital gains taxes. If at time 1 the investor sells the bond for a price above its book value ("tax basis"), she will realize a gain and pay tax at a rate τ_G. If the investor realizes a loss, conversely, she will receive a subsidy at the same rate τ_G. However, the subsidy applies only to the part of loss that is above the bond's face value.[27]

Define T_1 as the time-1 tax:

$$T_1 = -\tau_G \left(\max\left(1,\ P_1\right) - B_1\right) \tag{9}$$

where P_1 is the sale price at time 1, and B_1 is the tax basis at time 1. (A negative number means the investor pays a positive tax). For simplicity, B_1 is determined using straight line amortization:

$$B_1 = 1 + \frac{P_0 - 1}{2} = \frac{P_0 + 1}{2} \tag{10}$$

Also for simplicity, I assume that σ is large enough, so that $P_{\text{low}} < B_1$ and $P_{\text{high}} > B_1$. Finally, (9) contains the implicit assumption that premium bonds are optimal ($P_0^* \geq 1$). Later this assumption will turn out to be superfluous.

To clarify, consider the following three examples. Assume that the bond's face value is 100, but the bond was issued at a premium price of 110 at time 0. At time 1, the tax basis is $(110 + 100)/2 = 105$.

- If the bond is sold at 109, the *gain* is fully taxable and the investor pays a *tax* of $(109 - 105)\,\tau_G =$

[27] Based on the tax rules reviewed in Section A, the value of realizing losses above par is much larger than the value of tax trading under the quasi accrual, de minimis, and market discount regimes. To keep the math as simple as possible, I assume that the value of tax trading under all these other regimes is exactly zero.

$4\tau_G$.

- If the bond is sold at 101, the *loss* is fully taxable and the investor receives a *subsidy* of $(105-101)\,\tau_G = 4\tau_G$.
- If the bond is sold at 98, only *part of the loss* is taxable, because the price is now below face value. The investor receives a *subsidy* of $(105-100)\,\tau_G = 5\tau_G$ (instead of $(105-98)\,\tau_G = 7\tau_G$, if the loss were fully taxable).

This unusual set of rules is a stylized version of the actual U.S. federal tax code. Under the actual rules, the seller always get to deduct the full loss, as discussed in the previous Section. However, in the third example, the two dollars of discount to face value are called "market discount interest" and they will be taxable to the buyer as interest income sometime in the future. Because the buyer is likely a taxable entity, she will capitalize the future taxes into the price that she is willing to offer to the seller. This will make the price lower, negating part of the benefit of realizing losses for the seller.

D.2 The investor's problem

At time 1, when $P_1 = P_{\text{high}}$ (with probability p), the investor may be forced to sell the bond with an exogenously given probability λ. Moreover, when $P_1 = P_{\text{low}}$ (with probability $1-p$), the investor *will* sell and buy back the bond, because realizing losses is advantageous. The expected tax at time 1 is then

$$\mathbb{E}\left[T_1\right] = \tau_G \left[p\lambda \left(B_1 - P_{\text{high}}\right) + (1-p)\left(B_1 - \max\left(1,\ P_{\text{low}}\right)\right)\right] \tag{11}$$

And thus the investor will offer the following price for the bond:

$$P_0 = \delta_0 \left(\bar{\delta}\left(1+c\right) + c + \mathbb{E}\left[T_1\right]\right)$$

Use (10) and (11) to replace $\mathbb{E}\left[T_1\right]$ and solve for P_0:

$$P_0 = \delta_0 \frac{\bar{\delta}\left(1+c\right) + c + \tau_G L\left(c\right)}{1 - \delta_0 \tau_G \frac{p\lambda+1-p}{2}} \tag{12}$$

with

$$L\left(c\right) = \frac{p\lambda + 1 - p}{2} - p\lambda P_{\text{high}} - (1-p)\max\left(1,\ P_{\text{low}}\right).$$

D.3 The issuer's problem

The issuer is a state or local government, while the investor pays only Federal taxes. Thus, the issuer wants to design the security (i.e., choose c) in a way that minimizes the present value of expected taxes paid by the investor over the life of the bond. Because in practice the tax paid at time 1 will often turn out to be negative, I will frame the problem in terms of ***maximizing the tax subsidy***. It then becomes clear that what matters is not the dollar tax $\mathbb{E}[T_1]$, but rather the subsidy/proceeds ratio W:

$$\max_c W \equiv \frac{\delta_0 \mathbb{E}[T_1]}{P_0}$$

In other words, for every dollar of capital raised by issuing the bond, a fraction W comes from a federal tax subsidy, and does not entail any future cash outflows for the issuer.

Using again (10) and (11), substitute and simplify:

$$W = \delta_0 \tau_G \left[\frac{p\lambda + 1 - p}{2} + \frac{p\lambda \left(\frac{1}{2} - P_{\text{high}} \right) + (1-p) \left(\frac{1}{2} - \max(1,\ P_{\text{low}}) \right)}{P_0} \right] \tag{13}$$

Recognize that the numerator of the second term inside the brackets is just $L(c)$:

$$W = \delta_0 \tau_G \left[\frac{p\lambda + 1 - p}{2} + \frac{L(c)}{P_0} \right]$$

Use the definition of P_0 and simplify:

$$W = K + (1-K)\, \tau_G \frac{L(c)}{\bar{\delta}(1+c) + c + \tau_G L(c)}$$

$$K \equiv \delta_0 \tau_G \frac{p\lambda + 1 - p}{2} \in [0,\ 1]$$

Because K and $1 - K$ are positive constants,

$$\frac{\partial W}{\partial c} > 0 \iff \frac{\partial}{\partial c} \frac{L(c)}{\bar{\delta}(1+c) + c + \tau_G L(c)} > 0$$

which simplifies to

$$\left(c + \frac{\bar{\delta}}{1+\bar{\delta}} \right) L'(c) > L(c). \tag{14}$$

$L(c)$ contains a $\max$ operator, and thus it needs to be examined case by case.

D.3.1 Case I: $P_{\text{low}} < 1$

Intuitively, $P_{\text{low}} < 1$ should not be optimal, because only a part of the losses is tax deductible. This is easily verified by showing that the tax subsidy ratio is increasing in the coupon rate c. This would also imply that *some* premium is optimal, since par issuance ($P_0 = 1$) is a subcase of this case.

The definition of $L(c)$ in this case is

$$L(c) = \frac{p\lambda + 1 - p}{2} - p\lambda\left(\bar{\delta} + \sigma\right)(1+c) - (1-p)$$

$$L'(c) = -p\lambda\left(\bar{\delta} + \sigma\right)$$

Substituting the definitions of L, L' into (14) and simplifying, we obtain

$$\lambda\frac{p}{1-p}\frac{1-\bar{\delta}-2\sigma}{1+\bar{\delta}} < 1 \tag{15}$$

If $1 - \bar{\delta} - 2\sigma < 0$, W is always increasing in c. Otherwise, W is increasing in c under a very weak condition:[28]

$$\lambda\frac{p}{1-p} < \frac{1+\bar{\delta}}{1-\bar{\delta}-2\sigma} \tag{16}$$

The result that W is essentially always increasing in c makes things easier. So far we have assumed that issuing at a premium is optimal $(P_0^* \geq 1)$. A consequence of this result is that the assumption is essentially not needed, because

$$\left.\frac{\partial W}{\partial c}\right|_{P_0=1} > 0$$

and therefore issuing at par is suboptimal, and we can ignore all issue prices below par.[29]

[28] Essentially, this is an artifact of the straight-line book value accounting. When this condition is not verified, a capital gain at time 1 is likely to occur (high p) and to be realized (high λ); at the same time, increasing the coupon increases B_1 slower than P_{hi} and P_{lo}, so the capital gain increases while the capital loss decreases. In a binary tree model like the present one, any calibration not satisfying (16) would look contrived. Examples of calibrations that don't make the cut are ($\bar{\delta} = .7$, $\sigma = .05$, $\lambda = 1$, $p \geq 89\%$); ($\bar{\delta} = .1$, $\sigma = .01$, $\lambda = 1$, $p \geq 56\%$); etc.. For $p = 1/2$, (16) is always verified.

[29] Showing rigorously that below-par issuance is suboptimal would require specifying the OID tax rules mentioned in Section A, which would make the mathematics more complicated without adding any extra insight; especially given that the optimality of premium issuance should be obvious by just looking at Figure 13.

D.3.2 Case II: $P_{\text{low}} \geq 1$

After showing that some premium is optimal, the remaining question is whether there is a unique, well-defined optimal coupon. If

$$\left.\frac{\partial W}{\partial c}\right|_{P_{\text{low}} \geq 1} < 0$$

then the optimal coupon is the one that makes $P_{\text{low}} = 1$ exactly. In the opposite case, there is no unique optimum, and an infinite coupon is optimal.

An infinite coupon would be no problem in practice: all the arguments in this section are scale-invariant, because the issuer maximizes the *fraction* of issued principal that is a tax subsidy. For $c \to \infty$, this fraction is a well-defined number. An n-period bond with an infinite coupon is just like a mortgage with n payments, up to a scaling constant. Nonetheless, a closed-form, finite expression for the optimal coupon is desirable, because it makes it possible to make additional testable predictions.

The definition of $L(c)$ in this case is

$$L(c) = \frac{p\lambda + 1 - p}{2} - p\lambda\left(\bar{\delta} + \sigma\right)(1 + c) - (1 - p)\left(\bar{\delta} - \sigma\right)(1 + c)$$

$$L'(c) = -p\lambda\left(\bar{\delta} + \sigma\right) - (1 - p)\left(\bar{\delta} - \sigma\right)$$

Substituting the definitions of L, L' into (14) and simplifying, we obtain

$$-\frac{p\lambda}{1 - p}\left(1 - \bar{\delta} - 2\sigma\right) > 1 - \bar{\delta} + 2\sigma \tag{17}$$

This condition is hard, though not impossible, to meet: if $1 - \bar{\delta} - 2\sigma > 0$, the right-hand side is positive, but the left-hand side is negative.[30] If $1 - \bar{\delta} - 2\sigma < 0$, the expression simplifies to

$$\lambda\frac{p}{1 - p} > -1 - 4\frac{\sigma}{1 - \bar{\delta} - 2\sigma} \tag{18}$$

which can hold, but only in unrealistic parts of the parameter space.[31] Thus, a situation where

[30] This is likely to hold: with positive interest rates, $\bar{\delta} + \sigma < 1$ because it is a discount rate. If interest rates are more than one standard deviation away from zero, then $\bar{\delta} + 2\sigma < 1$, i.e. $1 - \bar{\delta} - 2\sigma > 0$.

[31] Again, with positive interest rates, $\bar{\delta} + \sigma < 1$ because it is a discount rate. Then, $1 - \bar{\delta} - \sigma > 0$, which implies that $\left|1 - \bar{\delta} - 2\sigma\right| < \sigma$, and thus the right-hand side is at least three. Even if $\lambda = 1$ (its highest possible value, corresponding to an investor who is forced to realize gains and losses at every period), one needs $p/(1-p) > 3$, i.e. $p > 3/4$. Like in the case of (16), in a binary tree model like the present one any calibration satisfying (18) would look contrived. Examples of calibrations that make the cut are ($\bar{\delta} = .5$, $\sigma = .49$, $\lambda = 1$, $p \geq 76\%$); ($\bar{\delta} = .7$, $\sigma = .29$, $\lambda = .5$, $p \geq 86\%$); etc. For $p = 1/2$, (18) is never verified.

an infinite coupon is optimal is possible (though unlikely) along with one in which there is a well-defined unique optimum.

D.4 Solution and comparative statics

When it exists (i.e. most of the time), the optimal coupon is found by setting $P_{\text{low}} = 1$, a "kink" solution like the one pictured in Figure 4.

$$P_{\text{low}} = \left(\bar{\delta} - \sigma\right)(1 + c^*) = 1$$

$$c^* = \frac{1 - \bar{\delta} + \sigma}{\bar{\delta} - \sigma} \tag{19}$$

The intuition is simple: the price should never go below par. If $P_{\text{low}} < 1$, when there are losses, some losses will not be deductible. Thus, increasing c is optimal because increasing c increases P_{low} and with it, the amount of losses that can be deducted. When $P_{\text{low}} > 1$, all losses are already deductible. As with all bonds, increasing the coupon reduces the duration of the bond because a larger share of the bond's cash flow is received at time 1 when the first coupon is paid. Thus, price volatility decreases, and with it decreases the value of the option to realize losses. The optimal price is

$$P_0^* = \frac{1 - \tau_G \bar{\delta} \frac{p\lambda + (1-p)}{2} + \sigma \left(1 - \tau_G \frac{(1-p) - 3p\lambda}{2}\right)}{\left(\bar{\delta} - \sigma\right)\left(\delta_0^{-1} - \tau_G \frac{p\lambda + 1 - p}{2}\right)} \tag{20}$$

Finally, rewrite by replacing the next period expected discount factor ($\bar{\delta}$) with the current discount factor plus expected change ($\delta_0 + \Delta$):

$$P_0^* = \frac{1 - \tau_G \left(\delta_0 + \Delta\right) \frac{p\lambda + (1-p)}{2} + \sigma \left(1 - \tau_G \frac{(1-p) - 3p\lambda}{2}\right)}{\left(\delta_0 + \Delta - \sigma\right)\left(\delta_0^{-1} - \tau_G \frac{p\lambda + 1 - p}{2}\right)}. \tag{21}$$

This expression will facilitate the interpretation of the dynamic programming results in the next Section.