Preliminary and Incomplete
Please Do Not Quote
Comments Welcome

The Price of a Tax*

Thomas J. Brennan†

February 28, 2014

Abstract

I define the price of a tax on an asset to be the price of the government's contingent claim on future cash flows, and I calculate it using risk-neutral pricing techniques. Levying an up-front lump-sum tax in this amount in lieu of an ex-post tax is equivalent to simply levying the ex-post tax, because the taxpayer and the government can invest in complementary self-financing dynamic replicating portfolios that allow markets to clear and preserve general equilibrium by reestablishing the same cash flows that would have occurred under the ex-post tax.

The price of a tax is a practical tool for quantifying the burdens of common non-linear and complex taxes. As an application, I analyze a tax that is linear for gains but disallows losses. With both numerical examples and theoretical results, I quantify how such a convex tax burdens risk-taking, particularly in the case of levered instruments like options. I also address how such a sub-additive tax burdens division of asset ownership using put-call parity or types of debt financing.

Keywords: Income Taxation, Proportionate Income Taxation, Loss Offsets, Progressivity, Graduated Rates, Risky Return to Assets, Domar-Musgrave, General Equilibrium
JEL Classification: D50, H20, K34

---

*I thank Alan Auerbach, Cathryn Brennan, Charlotte Crane, Dhammika Dharmapala, Ezra Friedman, Bill Gentry, Louis Kaplow, Bob McDonald, Dan Shaviro, Nancy Staudt, and Yair Listokin for many helpful comments and suggestions. The author also thanks participants at presentations of prior versions of this paper at the Yale Law, Economics & Organization Workshop in February 2010, the National Tax Association Annual Meeting in November 2009, and the NYU Tax Policy Colloquium in April 2009. Research support the Northwestern University School of Law Faculty Research Program is gratefully acknowledged.

†Professor of Law, Northwestern University School of Law and Professor of Finance (Courtesy), Northwestern University Kellogg School of Management, Chicago, IL 60611. Phone: 312-503-3233. E-mail: `t-brennan@law.northwestern.edu`

Contents

1 Introduction

It is well understood in the law and public finance literatures that a constant-rate tax levied proportionately on risky investment income is equivalent to a tax on the risk-free return to initial investment wealth.[1] What is less well understood is how precisely this result changes when the tax is not linear, which includes taxes with such commonly observed features as loss disallowances and graduated rates. This paper addresses the gap in understanding by extending the general equilibrium approach developed by Kaplow (1991, 1994) to a continuous-time setting and using asset pricing techniques for contingent claims in order to analyze non-linear methods of taxation.

One can think of a tax on future asset returns as a contingent claim held by the government. A taxpayer who ostensibly purchases the asset does not have the right to all future cash flows from it, but rather just to the residual amounts after the government's contingent claims are paid. The taxpayer therefore has a contingent claim as well, and this claim is complementary to that of the government, with the two claims together accounting for all future cash flows from the asset. I define the "price of a tax" to be the up-front price of the government's contingent claim, i.e., the amount the government would need to pay up-front in order to purchase its contingent claim.

To determine the price of a tax, it is useful to import tools from the financial economic theory of asset pricing. Specifically, assume that the price process for the risky asset is sufficiently well behaved that the government's contingent tax claim can be replicated by a self-financing dynamic portfolio of positions in the risky asset and a risk-free asset.[2] This will be the case case, for example, if the risky asset price follows a log-normal process.[3] The replicating portfolio has an initial cost, equal to the net cost of the initial risky and risk-free asset positions in the portfolio, and this is exactly the price of the contingent claim, i.e., the price of the tax.

If the government collects the price of the tax determined in this way, then it

[1]Domar and Musgrave (1944, 1945) lay out the result in its original form, Kaplow (1991, 1994) extends it to a general equilibrium setting, and Weisbach (2004a) extends this general equilibrium approach to the case of linear taxes levied at different rates on different assets. Further discussion and development in the legal literature can be found, for example, in Warren (1996) and Weisbach (2004b). There has also been critical treatment in the literature. See, for example, Avi-Yonah (2004) for a discussion of limitations on the result arising from borrowing costs and transaction costs.

[2]In order for this dynamic replication to work, the cash flows for the contingent claim must be deterministic at each point in time, conditional upon a particular state of the world for the asset value. However, the taxpayer may have discretion to choose, for example, the timing of tax cash payments and thereby affect the nature of the government's contingent claim. This taxpayer discretion will not affect the replicability of the claim, as long as the taxpayer's choices are deterministic, conditional upon a particular state of the world at a particular time. Under this assumption, the taxpayer's conditionally deterministic discretion can simply be incorporated into the conditionally deterministic rules governing the definition of the government's contingent claim.

[3]An example in which the desired type of dynamic replication would not be possible is a price process with stochastic volatility. In such a situation, dynamic replication would necessarily involve further instruments relating to the more complex volatility.

can manage a dynamically evolving self-financing investment portfolio to obtain all the cash flows it would have received under the ex post tax. The taxpayer is left with the initial amount he would have invested in the asset, reduced by the up-front tax collected by the government. This remaining amount is exactly the price of the replicating portfolio for the taxpayer's complementary contingent claim, because the two complementary contingent claims combined reflect all cash flows from the asset and thus the sum of their initial prices is simply the initial asset price. As a result, the government and the taxpayer can start with the price of the tax and the residual amount, respectively, and manage dynamic self-financing portfolios that will replicate the cash flows both parties would have had under the ex-post tax.[4] By doing this, the parties will ensure that markets clear, because their positions are complementary, and also that a state of general equilibrium will be preserved, assuming there was an equilibrium under the ex-post tax.

In a sense, all of this is simply a matter of relabeling, and nothing has really changed. The situation under the ex post tax is precisely replicated under the alternative up-front tax. The point of the exercise, however, is to determine the price of the tax. This value indicates how burdensome the tax on the asset is to the taxpayer, and it provides a present-value measurement tool that can be used to compare the relative burdens of taxes on different assets. It is also useful for government budgeting purposes because it represents the present value of the future tax collection rights.

The determination of the price of a tax can be complex, but this is necessary in light of the complex non-linear taxes applicable to investment returns. Currently, long-term capital gains are subject to graduated rates of 0%, 15%, or 20%, depending upon taxpayer income level,[5] and there is a 3.8% surtax on net investment income for taxpayers with sufficiently high incomes.[6] There is also a yearly limitation on net capital losses of $3,000 in per year for individuals.[7] In addition, investments in traditional IRAs are subject to taxation under progressive ordinary income tax brackets on withdrawal. Because of such non-linearities in the tax code, it is crucial to have adequate tools for understanding the burdens being imposed. The price of a tax is precisely the right type of tool for this job. It is theoretically grounded and also is extremely practical and effective in dealing with the complexities of the current tax rules.

My methodology creates a bridge between the pricing of contingent claims and the pricing of taxes in such a way that the many known theoretical and practical results regarding the former give rise directly to corresponding results for latter. Many of

[4]No further tax by the government should be levied on the taxpayer's dynamic replicating portfolio or transactions therein because the up-front payment amount represents full satisfaction of the taxpayer's obligations to the government with respect to the asset. The cash flows to the taxpayer from his replicating portfolio accordingly leave the taxpayer in precisely the after-tax cash flow position he would have been in under the ex-post tax.

[5]IRC § 1(h).

[6]IRC § 1411.

[7]IRC § 1211.

the most basic results for the pricing of contingent claims assume that markets are sufficiently complete that the claims may be dynamically replicated with portfolios of available assets, and in this paper I focus on this situation. I also make the further assumption that transaction costs are negligible. It is possible to relax these assumptions to significant degrees to obtain more general results, either finding exact prices or at least informative upper and lower estimates for prices.

In prior tax policy scholarship, the Domar-Musgrave result has been used as a strong premise for arguments that returns to capital should not be taxed, since taxation of such returns under an ideal linear tax is largely illusory and only the risk-free return to capital is actually effectively taxed. This paper takes a different approach by starting with the idea that the tax system has certain non-linear aspects, and that it is not feasible to eliminate these features completely. The methodology I develop provides a way to quantify the precise nature of the burdens imposed by such an imperfect system on taxpayers with different portfolios of assets, and this in turn informs tax policy by identifying what feasible changes are optimal to pursue. Such changes may include, for example, modifications in rates on certain assets or broader consideration of overall portfolio positions when determining tax due. This type of change may be second-best to a complete elimination of the tax on returns to capital, but it can still provide significant improvements in efficiency and substantial mitigation of some of the undesirable or unintended burdens imposed by non-linearity. The goal of my methodology is thus the provision of a robust set of tools for assessing critically the nature of an imperfect tax system and for guiding the development of incremental agendas for reform.

The price of a tax as I have discussed it thus far is an entirely general concept and can be applied to taxes that are due to various extents at multiple times in the future. In order to keep the analysis that follows more manageable, I generally focus on the special situation in which the ex post tax is only levied at a fixed time in the future. Similarly, I assume that after-tax cash flows to the taxpayer under the ex post tax also only occur at that same fixed time. These simplifying assumptions serve to make the details of the analysis easier, but they still provide a rich enough environment for the key ideas to be laid out.

To make the general theory more concrete, I describe in Section 3 the special case in which asset prices evolve over a finite number of periods and follow a simple binomial process. This model extends directly to an asset that follows a lognormal price process in continuous time.

As a particular example of a non-linear tax, I consider a tax that is linear for gains but does not allow offsets for losses. I denote this tax by T_{NL}. It is a convenient example for study because it is only slightly more complex than a linear tax, and because many real-world taxes do in fact deny offsets for some or all losses. To analyze the T_{NL} tax, I compute the price of a tax numerically and illustrate results graphically for various investment choices based on historical asset data. I also generalize beyond numerical results by proving theoretical propositions that confirm the properties shown by the calculations. Whenever possible, I prove these propositions not just for the T_{NL} tax, but also for broad classes of taxes that have certain similar properties that drive the results.

I find that convex taxes, such as the T_{NL} tax, generally burden the risk in risky returns, but not the expected level of these returns. As a result, such taxes encourage portfolio diversification targeted at risk minimization without regard to expected returns. This produces a distortion in investment allocation decisions. In addition, if a taxpayer owns put or call options on assets instead of the assets themselves, the T_{NL} tax can impose a particularly significant burden, which is larger for more out-of-the-money, and hence more levered, options.

I further find that subadditive taxes, such as the T_{NL} tax, penalize synthetic division of risky asset ownership across taxpayers. In other words, a lower aggregate burden is imposed on two investors who hold direct interests in a risky asset than would be imposed if the two people kept the same aggregate financial stake but arranged their affairs so that one would have a capped return (the debt claimant) and the other would have the upside beyond that return (the equity claimant).[8]

In terms of tax policy guidance, my results give preliminary support to the ideas that it may be desirable to have lower tax rates for taxpayers with relatively riskier portfolios and also to mitigate the aggregate tax burden on assets with debt-financed or synthetically divided ownership.

The remainder of the paper proceeds as follows. In Section 2, I provide a survey of some relevant related scholarship. In Section 3, I define notation and terminology and describe the finite binomial model and its continuous extension. In Section 4, I recover the classic Domar-Musgrave result as a special case of the general theory. In Section 5, I analyze the T_{NL} tax in detail. In Section 6, I address ways in which my methodology can be extended, as well as some limitations. Finally, Section 7 concludes.

2 Related Literature

The fact that a non-linear income tax may affect risky investment is not new. It is well-known that a tax disallowing losses, for example, preferentially favors gains over losses and generally distorts taxpayer behavior from what it would be if there were a linear tax, or no tax at all. Moreover, it is understood that any non-linearity in a tax system, such as a graduated rate brackets, can impose a significant burden on risk-taking.

Stiglitz (1969, 1972) is pioneering theoretical work demonstrating the burden of non-linear taxes on risk-taking. Atkinson and Stiglitz (1980) and Sandmo (1985)

[8]For a tax system that applies different rates and rules to debt and equity, debt financing is sometimes penalized and sometimes preferred. Which occurs depends on the specific nature of the asset, as well as the applicable rates and rules. This is discussed further in Section 5.4.

provide excellent reviews and analyses of this line of work, and Campisano and Romano (1981) apply this approach in the in their proposal for current recoupment of tax losses. The later work of Kaplow (1991, 1994) introduces a powerful general equilibrium methodology that does not depend upon expected utility maximization and innovates by introducing government investment portfolio adjustments as a tool for preserving equivalence between tax regimes. Weisbach (2004b) extends Kaplow's methods to the case of a linear tax imposed at different rates on different assets. In addition, the works of Schenk (2000) and Zelenak (2006) recognize that the usual arguments about the effect of linear taxes do not apply to non-linear taxes. Beyond the theoretical results and discussion, there is also an empirical literature, and the work of Gentry and Hubbard (2005) shows striking empirical evidence that entry into entrepreneurial activity by taxpayers is sensitive to convexity in the progressive rate schedule, with such convexity serving to discourage risk taking.

This paper moves beyond the existing literature by applying the general equilibrium techniques developed by Kaplow (1991, 1994) to the case of non-linear taxes. This innovation is accomplished by extending the existing model to allow for multiple periods and even continuous time trading. A further key contribution of this paper is the introduction of techniques from financial economics to evaluate the price price of the government's contingent tax claims on risky assets. This is what I term the price of a tax, and it has the property that an up-front lump-sum tax in its amount would be an equivalent alternative to the relevant non-linear tax. The price of a tax can be compared across investment strategies and quantifies the degree to which a tax burdens risk-taking of different types.

The results of this paper are also related to work in the corporate finance literature, where it has been long recognized that understanding of taxes on corporate income can be informed by the option pricing theory of financial economics. Green and Talmor (1985) and Majd and Myers (1986) present important research in this regard. The current paper is distinct from the corporate finance literature in that it employs the general equilibrium approach developed by Kaplow (1991, 1994) to analyze taxes in a context that takes into account fully the behavior of all parties, including the government. In addition, the current paper focusses on the question of the burden of a tax on individual investors instead of corporate investors. Nonetheless, there is significant connection between the corporate finance literature and the methodology set forth in this paper, and it is hoped that the two sets of literature may be able to benefit from and inform each other in the future.

3 Model with All Cash Flows at a Fixed Time

I extend the general equilibrium approach pioneered by Kaplow (1991, 1994) in two ways. First, I permit the tax levied to be non-linear. Second, I allow trading and rebalancing of investment portfolios, but no consumption of investment amounts, at points in between the start time, at which initial investments are made, and the end time, at which taxes are paid and after-tax investments are consumed. The idea is to import the ideas from asset pricing theory into this generalized model to determine the price for the value of the government-held contingent claim to collect taxes. The price of this claim truly represents the "price" of the tax, because imposing an up-front lump-sum tax in this amount, along with appropriate dynamic investment portfolio modifications, preserves a state of general equilibrium.

I focus most of the description on the specialized case of a "binomial" model that only allows a discrete number of asset prices at a discrete number of times. This treatment closely follows the ideas laid out in Cox et al. (1979), and it provides a rich framework that is relatively straightforward and intuitive. After discussing the binomial approach in detail, I explain generally what happens in the limit as the discrete model becomes continuous and how the price of a contingent claim may be thought of as the present value of an expectation taken with respect to an appropriate distribution.

Let $t_0 = 0$ and t_1 be the notation for the start and end times for investment, respectively. Also, let $n-1$ be the total number of evenly spaced intermediate times at which trading is allowed, and write these times as $t_{\frac{k}{n}} = \left(\frac{k}{n}\right) t_1$. Investors are given an endowment of wealth that they must invest at time t_0, and they must pay tax on their investment earnings according to a specified rule at time t_1. All after-tax wealth is consumed at time t_1, but no consumption occurs at intermediate points. The government may also act as an investor, but it is of course not subject to taxation. For the moment, all investments must be made in either a risky asset A or a riskless asset B. The return on A is uncertain, but B provides a guaranteed rate of return r, expressed on an annualized basis with continuous compounding. I write A_t and B_t to denote the price at time t for one unit of the assets A and B, respectively.

I start with the case $n = 1$, and I make the simplifying "binomial" assumption that A_{t_1} may take on only two possible values, namely uA_0 and dA_0, where $u > d$. In equilibrium, a taxable investor makes certain investment choices and is subject to tax in each of the two possible states of the world, corresponding to u and d. I write these two tax amounts as T_u and T_d. Because there are only two possible amounts of tax, and because the asset A takes on a different value in each of these states, it is possible to replicate the tax due to the government with the pre-tax amount that would come about if one started with a particular portfolio of assets. I determine the appropriate portfolio weights by solving the relevant system of two linear equations in two unknowns, and I find that the form of the portfolio is

$$P = \Delta A + sB, \quad \text{with} \quad \Delta = \frac{T_u - T_d}{(u-d)A_0} \quad \text{and} \quad s = \frac{uT_d - dT_u}{(u-d)B_{t_1}}. \tag{1}$$

I write P_t for the value of this portfolio at time t. If the government levies an amount P_0 of tax on the investor at the start time instead of levying any tax at the end time, and if the government invests this amount in the portfolio P, then the government will have the same cash flow at the end time as it would have had under the original tax. Also, if the investor, who now has P_0 less at the outset, invests in the portfolio $-P$ in addition to his otherwise planned investments, then his final cash flow will also be the same as it would have been under the original tax, because now no final tax is levied and the amount by which the final value of the $-P$ portfolio reduces his final wealth is exactly the same amount as the original tax did. Thus, replacing the original tax with this up-front tax and appropriate portfolio modifications preserves a state of general equilibrium. Consumption cash flows for all parties remain unchanged, and the offsetting adjustments of the taxpayer and the government insure that markets clear.

The analysis so far has been very specialized in that it is limited to the case $n = 1$ and permits only two possible values for A_{t_1}. Nevertheless, it demonstrates the powerful idea from asset pricing theory that it may be possible to replicate a contingent claim using a portfolio of the assets underlying the claim. In the simple case just considered, this is not very surprising, since the number of possible values for the government's contingent tax claim is the same as the number of possible values for the risky asset, and thus replication is merely a matter of solving a system of equations. Remarkably, however, the work of Cox et al. (1979) shows that is possible to extend the result to larger values of n and ultimately to pass to a continuous limit and generalize to a wide class of asset behaviors.

In order to handle the case $n > 1$, it is necessary to define what values of A are permitted at each point in time. The idea is to require that for a given value of A at a particular time, there are only two possible values of A at the next point in time, and these are multiplications of the original value by either u or d, where $u > d$. Thus, at time $t_{\frac{k}{n}}$, for $0 \leq k \leq n$, there are $k + 1$ possible values for A, namely $u^j d^{k-j} A_0$, for $0 \leq j \leq k$. Because A is the only uncertain quantity in this model, the values of A at the various times identify the possible "states of the world," and the permitted values for A make it clear that there are $k+1$ states of the world at time $t_{\frac{k}{n}}$. The tax due on the investment strategy of a particular taxpayer is completely specified by the amount of the tax in each of the $n + 1$ possible final states of the world, and I write the tax payable in these states as $T_{n,k}$, for $0 \leq k \leq n$. I claim that this specification of $n+1$ final tax values can be identified with a single value at time 0 , and that this number is the price of the tax. To see this, note that at each of the n states of the world at time $t_{\frac{n-1}{n}}$, there are only two possible values for the tax in the next period. Thus, just as in the $n = 1$ case described above, there is a portfolio that can be purchased in a particular state of the world at time $t_{\frac{n-1}{n}}$ that will produce the same pre-tax values as the amounts of the tax in the two possible future states. Thus, the $n + 1$ values for the tax at time t_1 give rise to n replicating portfolios at time $t_{\frac{n-1}{n}}$, and I write the prices of these replicating portfolios as $T_{n-1,k}$, for $0 \leq k \leq n - 1$. It is possible to continue this process repeatedly backward through time, with the next step being the calculation of $n-1$ prices labeled $T_{n-2,k}$, for $0 \leq k \leq n-2$, and so on, until a final value T_0 is obtained. At each time j, for $0 \leq j \leq n$, and for each state of the world k at that time, for $0 \leq k \leq j$, there is a particular replicating portfolio that the foregoing procedure identified, and it has value $T_{j,k}$.

If the government levies an up-front tax in the amount T_0, and then invests in replicating portfolios corresponding to each successive time and each state of the world that comes about, then the final result of this investment strategy is a value equal to the amounts of the original tax in each possible final state of the world, namely, $T_{n,k}$, for $0 \leq k \leq n$. Such successive investment is possible because the value of each portfolio at the time immediately after which it is purchased is exactly equal to the value of the next portfolio, by construction. If the government adopts this strategy in place of levying the original tax, then the taxpayer has T_0 less at time t_0, but he may adjust his investments in ways that exactly offset the government's portfolio strategy. Thus, as in the $n = 1$ case, replacing the original tax with this up-front tax and appropriate portfolio modifications preserves a state of general equilibrium, because consumption cash flows for all parties remain unchanged, and the offsetting adjustments of the taxpayer and the government insure that markets clear.

The portfolio adjustments required to transform a final tax into an up-front lump-sum looks like an elaborate game of what Kaplow (1994) termed "musical shares," and to a certain extent it is. The game pays significant dividends, however, inasmuch as it translates a complicated tax into a present value form that is amenable to ready comparison with the similarly calculated present values of other taxes. This "price" of the tax thus provides a common yardstick for evaluating the burdens imposed by any manner of tax.

One surprising feature of the binomial model I have presented so far is that there is no reference at all to the probability that the value of A moves up or down at any point. In fact, the results hold no matter what probability is used, and there was no need to make any particular choice. When pricing options in practice, a common assumption is that the probability that the price of A moves up, instead of down, at any discrete time step is given by a number q that is the same throughout the entire model. Cox et al. (1979) show that if

$$u = e^{\sigma\sqrt{t_1/n}}, \quad d = 1/u, \quad \text{and} \quad q = \frac{1}{2} + \frac{1}{2}\left(\frac{\mu}{\sigma}\right)\sqrt{t_1/n}, \tag{2}$$

then, as $n \to \infty$, the distribution of possible values for A tends toward a lognormal distribution with mean μ and standard deviation σ for the underlying normal distribution. This is the type of distribution assumed by Black and Scholes (1973) and is the assumption of the well-known Black-Scholes formula. In the continuous limit, with the probability and parameter assumptions, of (2), the binomial model used for option pricing leads to the result of Black and Scholes. Different choices for the parameters and probabilities in the binomial model can correspond to assets with non-lognormal behavior. For example, Cox et al. (1979) show how to incorporate jump diffusion processes as well.

In the continuous limit, the price of a contingent claim can be thought of as the initial value needed to fund a strategy of continuous dynamic portfolio investment in underlying assets, with the property that the chosen portfolio at any point in time precisely replicates the value of the contingent claim over an infinitesimal time interval. This powerful and innovative approach originated with Merton (1973, 1976), and he showed that this requirement leads to a stochastic differential equation which may be solved to determine the precise value of the claim at all points in time. The continuous-time approach works with my model in the same way that the discrete binomial model does. If the government levies an up-front tax equal to the price of its contingent tax claim, foregoes the collection of the original tax, and makes the appropriate continuous and dynamic investment choices, it can replicate the final cash flows it had under the original tax. The taxpayer has an initial value that is less by exactly the amount of the claim, and making portfolio adjustments that exactly offset the government's choices, he also replicates the final cash flows he had under the original tax. General equilibrium is again preserved, because all consumption cash flows are the unchanged, and markets clear at all times.

An additional way to view the value of a contingent claim is as the present value of the expected cash flows of the claim, where the expectation is taken with respect to a suitable probability measure for the distribution of asset values. Harrison and Kreps (1979) explain how this method for pricing claims encompasses those discussed above and is also based upon the same type of replicating portfolio argument. The probability measure most commonly used is the one that assumes all assets have an expected return equal to that of the riskless asset at all points in time. This is appropriate if there are risk-neutral agents in the economy that are either tax exempt or taxed continuously on a mark-to-market basis, because such agents will transact in such a way as to drive all expected returns to the same value. I will assume that my model includes such agents, so that this measure is correct, and with this choice of measure, discounting at the risk-free rate is appropriate. From this perspective, the price of a tax T that is a function of the difference between a stochastic final pre-tax investment wealth level, W_{t_1}, and an initial fixed wealth level, W_0, can be written

$$\mathcal{P}(A, T) = e^{-rt_1} \mathrm{E}\left[T\left(W_{t_1} - W_0\right)\right]. \tag{3}$$

The expectation E is taken with respect to the risk-neutral probability measure for final asset values, and r is the risk-fee rate of return, defined to be the rate of return on the riskless asset B, and assumed to be constant over time.

The expression for the price of a tax in (3) is very general, and it is the formulation that I will use throughout most of the remainder of this paper. It encompasses the possibility of multiple risky assets, as well as assets that follow distributions that may not be lognormal. For my purposes, it is not important to define precise boundaries for the scope of asset types that I consider, but I note that the formula in (3) holds as long as markets are sufficiently complete to allow the relevant dynamic replication portfolios to be formed at all points in time. Throughout the remainder of this paper, I assume that this type of market completeness holds. In Section 6, I discuss briefly extensions to situations of incomplete markets.

The price of a tax that I have defined truly indicates the burden of a tax inasmuch as it tells the amount of an equivalent up-front tax that can be levied in lieu of the original tax while still preserving general equilibrium. One caveat that is important to note, however, is that my methodology does not directly say anything about what happens if general equilibrium is altered, as may happen if tax rates or rules are changed. For example, if the tax rate applicable to a certain type of asset were increased, a new state of general equilibrium would occur, with taxpayers altering investment decisions and prices changing, both for the asset subject to the rate change and perhaps for other assets as well. In this new state of general equilibrium, a tax price for each investor could be computed, but it would depend not only upon the change in tax rules in isolation, but also on all of the concomitant changes to the rest of the economy. Of course it is possible to assume that the change in tax rules does not affect the state of general equilibrium, and then to analyze the effect of changes in tax rules on tax prices in the resulting partial equilibrium setting, and this approach may still provide valuable information notwithstanding the departure from the general equilibrium setting.

4 Revisiting the Linear Tax

The concept of the price of a tax defined in Section 3 can be applied to a linear tax to recover the well-known result of Domar and Musgrave (1944, 1945) and Kaplow (1991, 1994) that such a tax is equivalent to a tax on the risk-free rate of return to initial investment value. Under a linear tax, the government collects a constant fraction, τ, of returns to investment value. Thus, the payment from an individual to the government on an amount of income x is

$$T_{\mathrm{L}}(x) = \tau x. \tag{4}$$

It is possible for this payment to be negative, if $x < 0$, and in this case the government allows an offset of losses sustained.

Suppose that a taxpayer invests in a risky asset, A, for a period of length t_1. If A_t is the value of the asset at time t, then the amount the government is paid in taxes at time t_1 is equal to $\tau(A_{t_1} - A_0)$, including the possibility that this amount might be negative. The burden of the T_{L} tax on this investment is the price

$$\mathcal{P}(A, T_{\mathrm{L}}) = \tau e^{-rt_1}\mathrm{E}[A_{t_1} - A_0] = \tau e^{-rt_1}\left(A_0 e^{rt_1} - A_0\right) = \tau A_0\left(1 - e^{-rt_1}\right), \tag{5}$$

where the expectation is taken over the risk-neutral distribution of values for A_{t_1}, and so $\mathrm{E}[A_{t_1}] = A_0 e^{rt_1}$. This follows directly from the definition for the price of a tax in (3).

As an alternative to using the definition in (3) directly, the price of the T_{L} tax can be calculated by considering the prices of forward contracts and bonds. A forward is a contract that requires the purchase at time t_1 of an asset, A, for a specified amount. The purchase amount is called the forward price, and it is the number that causes the forward contract to have value zero at time $t = 0$. The forward price is thus given by $F_{t_1}(A) = e^{rt_1} A_0$, since this is the expected value of A_{t_1} under the risk-neutral distribution at time t_1, and so a contract that requires purchase of A for this price at time t_1 must have value zero at time zero. The claim of the government imposing the T_{L} tax on an investment in A can thus be seen to have a payoff identical to τ times the payoff of a forward contract on A at time t_1, plus τ times the gain on an investment of an initial amount A_0 in the riskless asset. The forward contract portion represents the proportional payment on the expected income over the risk-free return, which is zero under the risk-neutral expectation. The riskless investment portion represents the return on the risk-free return. This division of the T_{L} payoff into a forward and a riskless bond can be expressed symbolically as

$$T_{\mathrm{L}}(A_t - A_0) = \tau (A_t - A_0) = \tau \underbrace{\left(A_t - A_0 e^{rt_1}\right)}_{\text{Forward Payoff}} + \tau \underbrace{\left(A_0 e^{rt_1} - A_0\right)}_{\text{Bond Payoff}}. \tag{6}$$

The expected forward payoff is zero, since $\mathrm{E}[A_t] = A_0 e^{rt_1}$ under the risk-neutral expectation, and the expected bond payoff is $\tau A_0 (e^{rt_1} - 1)$. The sum of the present values of these two terms is $\tau A_0 e^{-rt_1}(e^{rt_1} - 1)$, which is of course the same as the price of the T_{L} tax obtained in (5).

It is informative to consider the details of the investment adjustments the government and a taxpayer would need to make to preserve general equilibrium if an up-front lump-sum tax in the amount of $\mathcal{P}(A, T_{\mathrm{L}})$ were imposed instead of the T_{L} tax. If the government invested the entire lump-sum amount in the riskless asset and also bought τ units of a zero-cost forward contract, it would obtain the same final payoff at time t_1 as it would under the T_{L} tax. If the taxpayer took the offsetting positions, borrowing the entire lump-sum amount at the risk-free rate and also selling τ units of a zero-cost forward contract, he would also obtain the same final payoff at time t_1 as he would under the T_{L} tax. If the parties preferred to transact in the asset directly, rather than in forwards, they could use the fact that a zero-cost forward contract to purchase an asset has exactly the same payoff as a long position in the asset, funded by riskless borrowing. Thus the τ units of long and short forward positions taken by the government and the taxpayer could be replaced by long and short positions in the asset, with the long position financed by borrowing at the risk-free rate and the proceeds from the sort position invested in the riskless asset.

5 Application To a Tax with No Loss Offsets

I next turn to an application of the general theory of Section 3 to a non-linear tax. I choose the example of a non-linear tax that is proportionate for gains but does not allow loss offsets. Thus, for an amount of income x, the tax levied is

$$T_{\mathrm{NL}}(x) = \tau \max(0, x). \tag{7}$$

This tax is of particular interest because it is not much more complicated than a linear tax, and yet it significantly burdens assets in ways that linear taxes do not. In addition, similar taxes are present in the current U.S. tax system, since net capital losses are disallowed or limited, while net capital gains are often taxed at a linear rate.[9] Study of the T_{NL} tax thus provides insight into the nature and magnitude of the burdens actually imposed by a non-linearity in the U.S. tax on capital gains.

The analysis of this section involves both concrete numerical examples and general theoretical propositions. For the numerical calculations, I make assumptions about underlying parameter values that are consistent with typical historical experience, and I display results graphically. The theoretical propositions generalize the numerical examples to statements that are independent of parameter values and other assumptions. The proofs of all propositions appear in the appendix.

5.1 Burden on a Single Risky Asset

To begin understanding the price of the T_{NL} tax, it is useful first to observe that the tax is convex, meaning that, for any two incomes, x and y, and for any weighting $0 \leq \alpha \leq 1$, the following inequality holds:

$$\text{Definition of Convexity:} \qquad T_{\mathrm{NL}}(\alpha x + (1-\alpha)y) \leq \alpha T_{\mathrm{NL}}(x) + (1-\alpha)T_{\mathrm{NL}}(y). \tag{8}$$

The proof that T_{NL} is in fact convex appears in Appendix A.1, and the fact that it is convex means that it burdens risk-taking in a sense made precise by the following proposition.

Proposition 1. *The price of a convex tax T with respect to an investment in a risky asset, A, is at least as great as the price of T with respect to an investment in the riskless asset, B. If the inequality defining the convexity of T is strict, then the price for A is strictly greater than the price for B. Also, if the inequality T satisfies is*

[9] An example of a loss limitation is the annual limit of \$3,000 under IRC §1211 on net capital losses for individuals. An example of a flat tax rate on gains is the 15% rate of IRC §1 that is generally applicable to long-term capital gains, other than those in certain categories, for taxpayers above a specified ordinary income tax bracket.

reversed, then T is said to be concave, and the price for A is no more than the price for B.

To describe the burden of T_{NL} on risk-taking in more detail, it is useful to write the price of this tax in terms of the value of a financial option. Specifically, if A is a risky asset that will be held until time t_1 and r is the annual risk-free rate of return, then the burden of T_{NL} on this investment is

$$\begin{aligned}\mathcal{P}(A, T_{\mathrm{NL}}) &= e^{-rt_1}\mathrm{E}\left[T_{\mathrm{NL}}\left(A_{t_1} - A_0\right)\right] \\ &= \tau e^{-rt_1}\mathrm{E}\left[\max\left(0, A_{t_1} - A_0\right)\right].\end{aligned}$$

The expectation in this equation is taken with respect to the risk-neutral distribution for A_{t_1}, and the present value of this expectation is exactly equal to the price of a European call option on A with strike A_0 and expiration at time t_1. The price of this call may be denoted as $C_{t_1}(A)$, and then the burden of the tax may be written $\mathcal{P}(A, T_{\mathrm{NL}}) = \tau C_{t_1}(A)$

The theory of option pricing may be applied directly to shed light on the nature of $\mathcal{P}(A, T_{\mathrm{NL}})$. In particular, if the price of A has a log-normal distribution with underlying volatility σ, then the well-known formula of Black and Scholes (1973) provides a value for the call option $C_{t_1}(A)$, and the price of the tax is

$$\mathcal{P}(A, T_{\mathrm{NL}}) = \tau A_0\left(N(d) - N\left(d - \sigma\sqrt{t_1}\right)e^{-rt_1}\right), \tag{9}$$

where $d = \left(\frac{\sigma}{2} + \frac{r}{\sigma}\right)\sqrt{t_1}$, where N is the cumulative distribution function for the standard normal distribution, so that $N(x) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x} e^{-x/2}\,dx$.

Proposition 2. *If A has a log-normal distribution with underlying volatility σ, then the burden of the T_{NL} tax on an investment in A of length t_1 can be expressed as*

$$\mathcal{P}(A, T_{\mathrm{NL}}) = \tau A_0\left(\frac{1}{\sqrt{2\pi}}\right)\left(\sigma\sqrt{t_1}\right) + E,$$

where $|E| \leq 5\tau A_0\sigma^2 t_1$, provided $r \leq \sigma^2$ and $\sigma^2 t_1 \leq 1$. Thus, the approximation $\mathcal{P}(A, T_{\mathrm{NL}}) = \tau A_0\left(\frac{1}{2\pi}\right)\left(\sigma\sqrt{t_1}\right)$ is roughly accurate for small values of $\sigma^2 t_1$.

Proposition 2 shows that, when $\sigma^2 t_1$ is not too large, the burden of the T_{NL} tax on a risky investment with log-normal returns is approximately proportionate to the risk, as measured by volatility. Thus, the T_{NL} tax can be seen to burden risk-taking in a very direct way. Figure 1 illustrates this burden for varying levels of volatility, σ, and for varying times of investment, t_1, with the underlying calculations carried out using the formula in (9). It is apparent from Figure 1(a) that for the fixed investment horizon $t_1 = 1$, the burden grows proportionately with volatility. Figure 1(b) shows that higher volatilities continue to have a higher tax burden as t_1 grows, but that the burden for any fixed volatility increases in time.

The parameter assumptions for the calculations in Figure 1 are based on historic values computed from data available at the website of Professor French.[10] These data underlie many of the numeric calculations in this section and are summarized in Table 1.

Parameter	Value	Description
r	3.6%	Risk-free rate, computed as 12 times the average of the monthly risk-free returns specified in Professor French's data for the period from July 1926 to May 2011
σ_{mkt}	18.9%	Volatility, computed as $\sqrt{12}$ times the standard deviation of monthly market returns specified in Professor French's data for the period from July 1926 to May 2011
$\mu_{\mathrm{mkt}} - r$	7.5%	Market excess return, computed as 12 times the average of the monthly market returns in excess of the risk-free rate in Professor French's data for the period from July 1926 to May 2011
$S_{\mathrm{mkt}} = \frac{\mu_{\mathrm{mkt}}-r}{\sigma_{\mathrm{mkt}}}$	0.40	Sharpe ratio, computed as the ratio of the historic $\mu_{\mathrm{mkt}} - r$ value to the historic σ_{mkt} value

Table 1: Historic quantities used as parameters for certain numerical calculations

Figure 1(b) suggests that the burden of the T_{NL} tax increases monotonically with the length of the investment. The next proposition provides a precise formulation of the size of the burden for large values of t_1. Although the proposition shows that the burden on any asset ultimately tends to the limit τA_0 as $t_1 \to \infty$, it is clear from Figure 1(b) that this limit is not reached, under reasonable parameter assumptions, for investment periods of up to 30 years.

Proposition 3. *If A_{t_1} has a log-normal distribution with underlying volatility σ, then the burden of T_{NL} on an investment in A of length t_1 satisfies*

$$\tau A_0 \left(1 - N\left(d - \sigma\sqrt{t_1}\right)\left(1 + e^{-rt_1}\right)\right) \leq \mathcal{P}\left(A, T_{\mathrm{NL}}\right) \leq \tau A_0,$$

and the left-hand side tends to the limit τA_0 as $t_1 \to \infty$. In addition, the burden of T_{NL} on an investment in the riskless asset B satisfies

$$\tau B_0 \left(1 - e^{-rt_1}\right) \leq \mathcal{P}\left(B, T_{\mathrm{NL}}\right) \leq \tau B_0.$$

Thus the percentage burden on A and B tends toward the same limit, τ, as $t_1 \to \infty$.

[10]The data were accessed on August 2, 2011 from Professor French's data library on his faculty web page at `http://mba.tuck.dartmouth.edu`.

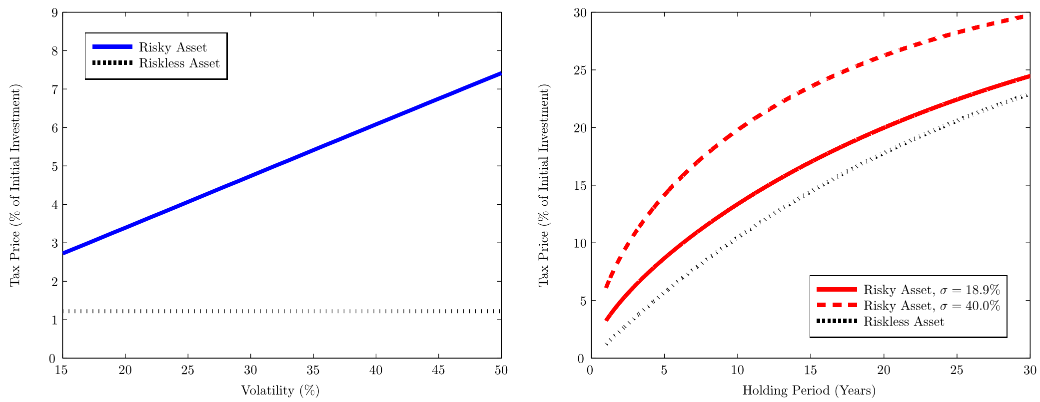

(a) Varying σ Values

(b) Varying t_1 Values

Figure 1: The figure on the left shows how the burden of the T_{NL} tax depends upon the volatility σ in the case of an investment of $t_1 = 1$ years in a risky asset with log-normally distributed values. The risk-free rate is assumed to be the historic average of $r = 3.6\%$, and the horizontal dotted line shows the burden of the T_{NL} tax on the riskless asset for comparison. The figure on the right shows how the burden of the T_{NL} tax changes with the length of investment, t_1, which varies along the horizontal asset. The risky asset is again log-normally distributed, and the assumed volatility is the historic market volatility $\sigma = 18.9\%$. The upper dashed curve indicates how the picture would change if the volatility were higher, with $\sigma = 40\%$, and the lower dotted curve indicates the burden on the riskless asset for comparison. For both graphs, the assumed tax rate on gains is $\tau = 35\%$.

Proposition 2 also shows that the burden of the T_{NL} tax depends upon asset return volatility, but not on the market expectation for asset return level. This is a direct consequence of the fact that the formula for the price of a tax involves an expectation with respect to risk-neutral outcomes, rather than market expected outcomes. Nevertheless, if there is a functional relationship between volatility and market expectations for asset return level, the T_{NL} tax places an indirect burden on asset returns by virtue of this relationship. For example, if a one-factor model for asset returns holds, such as is the case with the Capital Asset Pricing Model (CAPM), then

$$\mu_A = r + \beta_A \left(\mu_{\mathrm{mkt}} - r \right) \quad \text{and} \quad \sigma_A = \beta_A \sigma_{\mathrm{mkt}} + \sigma_\varepsilon,$$

where μ_A and μ_{mkt} are the market expectations for return levels for the asset A and the overall market, respectively, where σ_A and σ_{mkt} are the corresponding volatility numbers, where β_A is a constant dependent on the asset A, and where σ_ε is the volatility of the idiosyncratic component of the returns for asset A that is uncorrelated with market returns. As usual, r denotes the risk-free rate. These equations combine to show that

$$\sigma_A = \frac{1}{S_{\text{mkt}}}(\mu_A - r) + \sigma_\varepsilon,$$

where $S_{\text{mkt}} = \frac{\mu_{\text{mkt}} - r}{\sigma_{\text{mkt}}}$ is the Sharpe ratio for market returns. Proposition 2, along with the historic value of $S_{\text{mkt}} = 0.4$ from Table 1, therefore imply that[11]

$$\mathcal{P}(A, T_{\text{NL}}) \approx \tau A_0 (\mu_A - r)\sqrt{t_1} + \tau A_0 \sigma_\varepsilon \sqrt{t_1}. \tag{10}$$

Thus, based on assumptions compatible with the CAPM, the burden of the T_{NL} tax over a one-year period is equal to a sum of two components, the first of which is equal to τ times the market-expected return on A over the risk-free rate, and the second of which grows with the idiosyncratic risk of A. For longer time periods, the entire burden grows by a factor of the square root of the length of time. This finding stands in sharp contrast to the well-known result from Section 4 that linear taxes only burden the risk-free return to initial invested capital.

5.2 Portfolios of Multiple Assets

Because taxpayers often hold more than one asset simultaneously, it is desirable to analyze the burden of T_{NL} on portfolios containing multiple assets. In the simplest case, a taxpayer may split his initial wealth between a risky asset, A, and the riskless asset, B. Figure 2(a) illustrates how the tax burden depends on the initial allocation choice in the case of a one-year investment in such a portfolio. The parameters for the assets are those specified in Table 1, with A having the risk of the historic market portfolio and with the risk-free rate equal to the historic average. The most notable feature of Figure 2(a) is that although the burden of the T_{NL} tax generally increases with the size of the allocation to the risky asset, it does not do so linearly. Instead, for a relatively small allocation of initial wealth to the risky asset, the burden is nearly the same as it would be for an allocation entirely to the riskless asset. It is only above a threshold of about a 15% investment in the risky asset that the burden increases in a roughly linear fashion. This is a reflection of the fact that, for small allocations of wealth to the risky asset, potential losses from the risky position are generally outweighed by guaranteed gains in the riskless asset. As long as the net change in value of the portfolio over time is positive, the non-linear tax T_{NL} is the same as its linear counterpart, T_{L}, and this latter tax levies the same tax burden on initial investment wealth, regardless of investment choice. Thus, the T_{NL} tax does not discourage all risk-taking in these simple portfolios, but only that risk-taking that is

[11] Note that the approximate equality $S_{\text{mkt}}\sqrt{2\pi} = 0.4\sqrt{2\pi} \approx 1$ is used to obtain the approximate price formula.

creates a significant probability for net portfolio losses. This principle is formulated precisely in the following proposition.

Proposition 4. *If there is a positive probability that $A_{t_1} < A_0$, with respect to the risk-neutral distribution for A_{t_1}, then the burden of the T_{NL} tax on an investment in A for time t_1 is strictly greater than the burden of the corresponding linear tax. In the alternative, if $A_{t_1} \geq A_0$ has probability 1, then the burdens of the two taxes are identical.*

If a taxpayer holds a portfolio of more than two risky assets, the situation is more complex because there are several dimensions along which he may vary his portfolio decision. The optimal choices, however, fall along a single-dimensional "efficient frontier," which is defined by the property that each portfolio on the frontier has the highest possible expected return for its given level of risk. In the absence of taxes, an investor would generally choose an optimal portfolio along this frontier, with the particular choice dictated by his relevant utility function, assuming the utility function depends only upon the risk and return of the portfolio. If the T_{NL} tax is imposed, however, the tax price will generally be higher for higher for portfolios with greater risk, and so this tax tends to distort the optimal portfolio choice toward a lower risk alternative. In contrast, if the T_{L} tax were imposed, the selection of a portfolio would of course be independent of the tax, since the T_{L} tax is equivalent to a tax on initial wealth, regardless of investment choices.

To get a sense of the types of distortion that may occur if the T_{NL} tax is imposed on the returns of a portfolio of several assets, it is useful to calculate a numerical example. Consider three asset choices with expected returns and risks equal to the historic values for the market portfolio, the "small minus big" (SMB) portfolio, and the "high minus low" (HML) portfolio described on the website of Professor French.[12] I label these assets A, A^{SMB}, and A^{HML}, and I compute annualized historic returns and risks for A^{SMB} and A^{HML} based on monthly data from July 1926 to May 2011 in the same way I computed these values for A in Table 1. I also compute monthly correlations between each pair of assets using the same historic data. The annualized returns are 11.1%, 3.0%, and 4.7% for A, A^{SMB}, and A^{HML}, respectively, and the corresponding annualized risks are 18.9%, 11.5%, and 12.4%. The correlation between A and A^{SMB} is 0.33, that between A and A^{HML} is 0.23, and that between A^{SMB} and A^{HML} is 0.10. Figure 2(b) shows the efficient frontier calculated using these parameter values, and it highlights the portfolio with the maximum pre-tax return-to-risk ratio, as well as the portfolio with the minimum pre-tax risk. The price of the T_{NL} tax is lowest for the latter portfolio, but the former will be preferred in the absence of taxes by investors desiring to maximize their return-to-risk ratio.

The optimum choice for an investor is not necessarily the portfolio with the minimum tax price. Nonetheless, an investor who would otherwise choose a higher reward-to-risk ratio may satisfy himself with a somewhat lesser ratio in order to achieve a

[12] See footnote 10.

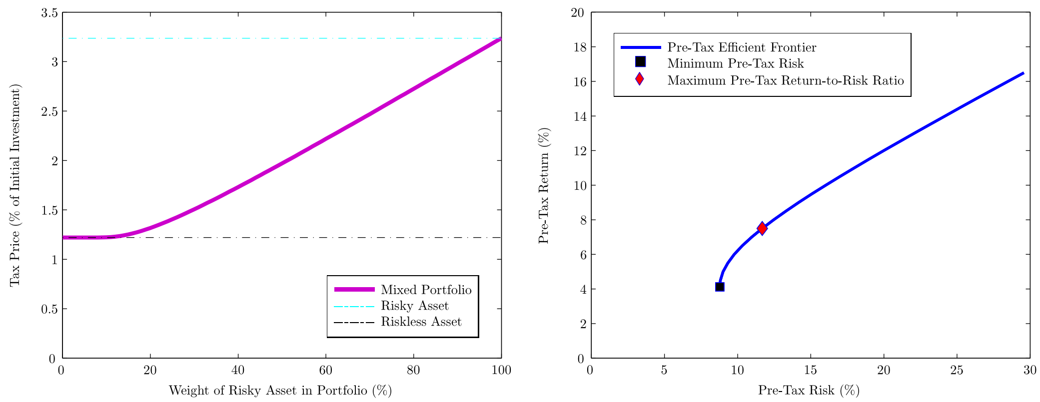

(a) Risky and Risk-free Asset

(b) Multiple Risky Assets

Figure 2: The figure on the left illustrates the price of the T_{NL} tax in the case of a portfolio that has a fraction of initial wealth invested in a risky asset with a volatility of 18.9% and the remaining wealth invested in the riskless asset, with a guaranteed return of 3.6%. The period of the investment is one year, and the tax rate on gains is $\tau = 35\%$. The fraction of investment in the risky asset varies from 0% to 100% along the horizontal axis, and the dashed horizontal lines represent the price of the tax on an investment in exclusively one of the two possible assets. The figure on the right illustrates the mean-variance efficient frontier for the three assets A, A^{HML} and A^{SMB} described in the text. It labels the portfolio of minimum pre-tax risk on the frontier, which is the portfolio for which the burden of the T_{NL} tax is smallest. It also labels the portfolio with the maximum pre-tax Sharpe ratio, for which the burden of the T_{NL} tax is significantly higher.

smaller tax burden, depending upon the relevant underlying utility function. To the extent investors respond to imposition of the T_{NL} tax in this way, portfolio allocation decisions is altered, and optimal allocation of capital is not achieved. The statistics summarized in Table 2, based on the numerical example involving A, A^{SMB}, and A^{HML}, show that the degree to which allocations differ may be substantial.

5.3 Analysis of Put and Call Options

In addition to analyzing traditional assets, it is also informative to consider the burden imposed by the T_{NL} tax on derivative instruments, such as European put and call options. Because these instruments represent leveraged positions implicitly, their returns tend to have greater volatility, and this generally corresponds to a higher price of the T_{NL} tax, relative to the price for an investment in the underlying asset.

Figure 3(a) illustrates the price of the T_{NL} tax for an investment exclusively in a one-year call option, $C(K)$. The asset, A, underlying the options is assumed to have the historic risk of the market, and the risk-free rate is assumed to be the historic

	Portfolio with Minimum Risk	Portfolio with Maximum Reward-to-Risk Ratio
Tax Cost (% of Wealth)	1.9%	2.3%
Risk	8.8%	11.7%
Return	4.1%	7.5%
Reward-to-Risk Ratio	0.47	0.64
Fraction in A	0.05	0.48
Fraction in A^{SMB}	0.51	0.13
Fraction in A^{HML}	0.44	0.39

Table 2: Description of pre-tax minimum risk and maximum reward-to-risk ratio portfolios

average, with both values as specified in Table 1. The strike price, K, of the call is expressed as a percentage of initial asset value and varies along the horizontal axis of the graph. To calculate the price of the tax, a taxpayer is assumed to choose a value of K and invest his entire initial wealth in $C(K)$. The price of the tax for a one-year investment of this type is reported for a range of possible values of K. The value of K varies along the horizontal axis and is expressed as a percentage of initial asset value. The price of the tax is expressed along the vertical axis as a percentage of initial wealth. Note that because $C(K)$ is less expensive with increasing strike, the investor is assumed to purchase a larger number of calls if he selects a higher value of K for his strategy. Nonetheless, the total wealth invested is the same across all taxpayer choices of K.

For low values of K, the payoff on $C(K)$ is closer to that on the underlying risky asset, and when $K = 0$, $C(0)$ is identical to the underlying risky asset. Accordingly, the price of the T_{NL} tax for a portfolio of call options tends to the price of the T_{NL} tax for investment in the underlying asset as the strike price tends toward zero. On the other hand, for high values of K, positive payoffs are rare, but generally large when they occur, relative to initial cost of $C(K)$. As a result, τ times the price of $C(K)$, which is the present value of the expected payoff, is close to the price of the tax for an investment in $C(K)$, which is τ times the present value of the expected profit, defined as the expected payoff less the initial cost. Accordingly, the price of the T_{NL} tax for a portfolio of call options tends to a fraction τ of the initial wealth as K grows large. Figure 3 illustrates these two extreme behaviors, with the price tending toward $\mathcal{P}(A, T_{\mathrm{NL}})$ as $K \to 0$ and toward a fraction τ of initial wealth as $K \to \infty$. These tendencies do not depend upon the particular parameters used, and Proposition 5 provides a theoretical generalization.

The case of a put option, $P(K)$, is complementary to that for the corresponding call, with the price of the T_{NL} tax being greater for lower values of the strike, K, rather than for higher ones. In addition, the price of the tax tends toward $\mathcal{P}(A, T_{\mathrm{NL}})$

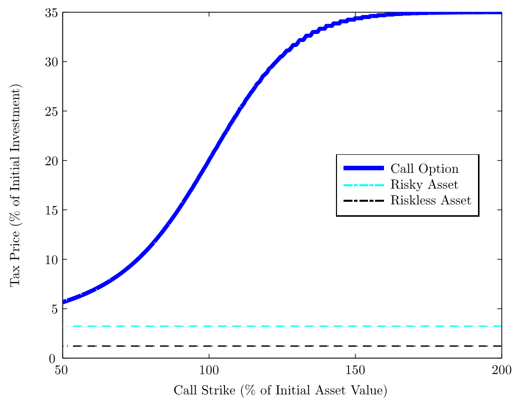

(a) Price of T_{NL} Tax for Call Option

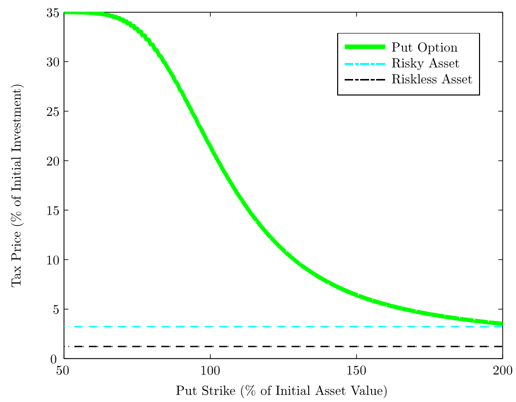

(b) Price of T_{NL} Tax for Put Option

Figure 3: The figure on the left shows the price of the T_{NL} tax for a taxpayer who invests entirely in call options on the risky asset, A, with strike price K. The value of K varies along the horizontal axis and is expressed as a percentage of A_0, the initial value of A, and the vertical axis indicates the price of the T_{NL} tax as a percentage of initial wealth invested. For comparison, the two dashed lines show the price of the T_{NL} tax for a taxpayer investing in either just A or just the riskless asset B. The figure on the right shows the analogous values for an investment in put options on A with strike price K.

as K grows large, and toward a fraction τ of initial wealth as K tends to zero. The intuition for each of these tendencies is analogous the corresponding results for call options described above. Proposition 5 provides a theoretical generalization in this case as well.

Proposition 5. *If a taxpayer invests a fixed amount of initial wealth, W, in call options with strike K, then the burden of the T_{NL} tax on this investment increases with K. At the extreme limits, it satisfies*

$$\lim_{K \to 0^+} \mathcal{P}\left(C(K), T_{\text{NL}}\right) = \mathcal{P}\left(A, T_{\text{NL}}\right) \quad \textit{and} \quad \lim_{K \to \infty} \mathcal{P}\left(C(K), T_{\text{NL}}\right) = \tau W.$$

If the investor instead pursues the same strategy for European puts instead of European calls, the tax burden under the T_{NL} tax decreases with the strike price of the options, and at the extreme limits, it satisfies

$$\lim_{K \to 0^+} \mathcal{P}\left(P(K), T_{\text{NL}}\right) = \tau W \quad \textit{and} \quad \lim_{K \to \infty} \mathcal{P}\left(P(K), T_{\text{NL}}\right) = \mathcal{P}\left(A, T_{\text{NL}}\right).$$

5.4 Put-Call Parity and Debt Financing

Debt financing for an asset can be thought of in terms of put and call options,[13] and an analysis of the price of the T_{NL} tax for appropriate option positions can thus provide an indication of the burden imposed by the tax on the use of such financing. The idea is to think of the equity owner as having a call option on the asset, with a certain strike price, K, and to think of the debt owner as having a short put position with respect to the asset, also with strike price K, as well as a position in the riskless asset that has a final payment of K, including both principal and interest. The combined positions of the debt owner constitute a "funded short put," because the position in the riskless asset provides a large enough guaranteed payment that the combined position will never result in a negative aggregate final payment owed by the debt owner. The combination of the equity position and the debt position results in complete ownership of the asset so that the following relationship holds:

$$\text{Put-Call Parity:} \qquad \text{Risky Asset} = \underbrace{\text{Call Option}}_{\text{Equity}} + \underbrace{\text{Funded Short Put}}_{\text{Debt}}.$$

The burden of the T_{NL} tax on debt financing may thus be thought of as the sum of the price of the tax for the equity owner and the price for the debt owner, or

$$\text{Burden of } T_{\text{NL}} \text{ Tax on Debt Financing} = \mathcal{P}\left(C(K), T_{\text{NL}}\right) + \mathcal{P}\left(FSP(K), T_{\text{NL}}\right), \quad (11)$$

where $C(K)$ is the call option, $FSP(K)$ is the funded short put, and K is the common strike price.

The burden of the T_{NL} tax on call options is analyzed in Section 5.3, and for funded short puts, sample calculations are illustrated in Figure 4(a). Because the funded short put behaves much like the riskless asset for low strike values and much like the underlying risky asset for high strike values, the tax price for a funded short put tends toward the tax prices for the riskless and risky assets as the strike becomes small or large, respectively. All of the calculations in Section 5.3, as well as those underlying Figure 4(a) assume that a constant amount of wealth is invested, meaning that more instruments were purchased when the value of a single instrument was low. In order to combine these results to arrive at the aggregate burden described in (11), it is necessary to rescale tax prices to reflect the fact that the initial wealth invested by both the equity and debt owners is equal to the initial cost of their respective positions. Figure 4(b) shows the result of the required rescaling and aggregation in the case of a risky asset with historic market volatility and a risk-free rate equal to

[13]The approach to thinking of debt financing dates back to Merton (1974), but while that paper specifically considers corporate debt, there is no assumption here that the debt be that of a corporation.

the historic average. For these parameters, it is appears that the burden of the T_{NL} tax on debt financing is generally greater than the tax burden on direct ownership of the risky asset. That is to say, debt financing is penalized by the T_{NL} tax.

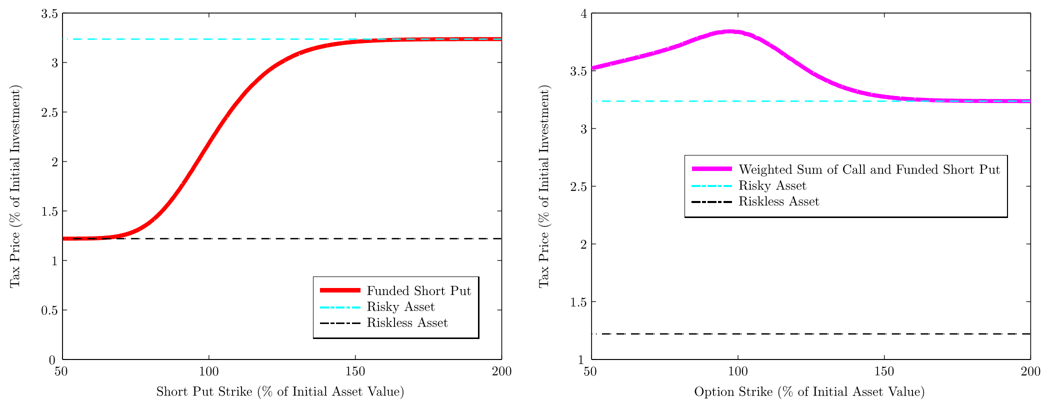

(a) Funded Short Put (b) Divided Ownership Using Put-Call Parity

Figure 4: The figure on the left shows the price of the T_{NL} tax for a one-year investment in a funded short put position. The figure on the right shows the aggregate price of the T_{NL} tax for two investors, one taking a funded short put position with respect to the same underlying asset, and the other taking a long call position with respect to the same asset and with the same strike price as the funded short put. All options considered are on an underlying asset with a volatility of 18.9%, the risk-free rate is assumed to be 3.6%, and the rate of tax on gains is $\tau = 35\%$. In both graphs, the strike price varies along the horizontal axis and is expressed as a percentage of the current value of the risky asset. Also, the dashed lines represent the price of the T_{NL} tax on an investment in either just the riskless asset or just the risky asset underlying the options. The price of the tax in all cases is expressed as a percentage of total initial wealth invested.

Figure 4(b) shows that debt financing can increase the burden of the T_{NL} tax up to 18.7% over the burden imposed on direct asset ownership. This highest amount of increase for the parameters used occurs when the strike price is $K = 96.7\%$ of the initial price of the underlying asset. For this strike price, the (pre-tax) cost of the funded short put is 88.9% of the initial asset price and that of the call is 11.1% of the initial asset price. The price of the T_{NL} tax for the equity owner is 18.4% of the (pre-tax) price of the call, and the price for the debt owner is 2.0% of the (pre-tax) price of the funded short put. As a result, the aggregate burden to both investors is

$$\underbrace{88.9\% \times 2.0\%}_{\text{Debt} \times \text{Tax Price of Debt}} \quad + \quad \underbrace{11.1\% \times 18.4\%}_{\text{Equity} \times \text{Tax Price of Equity}} \quad = \quad \underbrace{3.8\%.}_{\text{Aggegate Tax Price}}$$

This aggregate tax price is expressed as a percentage of initial asset value, and it is 18.7% higher than 3.2%, the price of the tax on an investment directly in the risky asset. These calculations show that not only is the aggregate burden of the T_{NL} tax on investment increased when debt financing is used, but also that the burden is spread unevenly across the two types of owners. In this example, the equity owners face a tax price that is nearly six times higher than that for direct asset ownership, while debt owners have a tax price that is less than that for direct asset ownership, but still significantly higher than for ownership of pure debt, namely the riskless asset.

The pattern of behavior illustrated in Figure 4(b) holds true more generally for other choices of parameters and for taxes other than just the T_{NL} tax. Proposition 6 states this result precisely for a subadditive tax, which defined to be a tax T, such that for any two incomes x and y, the following relationship holds:

$$\text{Subadditivity} : T(x+y) \leq T(x) + T(y). \tag{12}$$

Proposition 6. *Let A be a risky asset that has ownership divided between a debt holder and an equity owner. If T is a subadditive tax, then the price of T for direct ownership of A is less than or equal to the sum of the prices of the tax for the debt owner and the equity owner. That is, debt financing has at least as high a tax price as direct ownership. If T satisfies the reverse inequality of that in* (12)*, then it is said to be superadditive, and the tax burden associated with debt financing is less than or equal the tax burden on direct ownership.*

The following corollary provides an even stronger result in the particular case of the T_{NL}, rather than a generic subadditive tax.

Corollary 1. *If the distribution of possible final values for the risky asset has support throughout a neighborhood around the strike price K, then strict inequality holds in the result of Proposition 6 in the case of the T_{NL} tax. That is, the aggregate tax burden is strictly higher for debt financing than the corresponding burden for direct ownership.*

It is possible that debt and equity owners may be subject to different rates of tax, and perhaps even allowed loss offsets to differing degrees, and it is interesting to analyze how such features affect the aggregate tax price for debt financing. Figure 5 illustrates the results of calculations that assume that the debt owner is taxed at a rate of $\tau_D = 35\%$ and that the equity owner is taxed a preferential rate of $\tau_E = 15\%$. In addition, the calculations encompass both the situation in which loss offsets to the debt owner are disallowed, and that in which deductions for losses to the debt owner are permitted. The underlying risky asset has the historic volatility of the market, and the risk-free rate is the historic average. The horizontal axes in the figure represent the initial debt-equity ratio, measured as the ratio of the value of the debt to the value of the equity at time $t = 0$. For the range of initial debt-equity ratios up to 10 shown in Figure 5(a), there is a signficiant burden to debt financing beyond that for simple ownership of the underlying asset, although the burden is somewhat reduced for debt-equity ratios above 2 if the debt owner is permitted an offset for losses. For some of the much higher ratios illustrated in Figure 5(b), debt financing is seen to have a burden somewhat lower than direct asset ownership. For such high ratios, however, it is likely that the debt would not be treated as such for tax purposes, and both co-owners would actually receive equity treatment. In this case, the results illustrated in Figure 4, in which all debt and equity is taxed at the same rate, provide a more accurate description of what the relevant tax burden actually is.

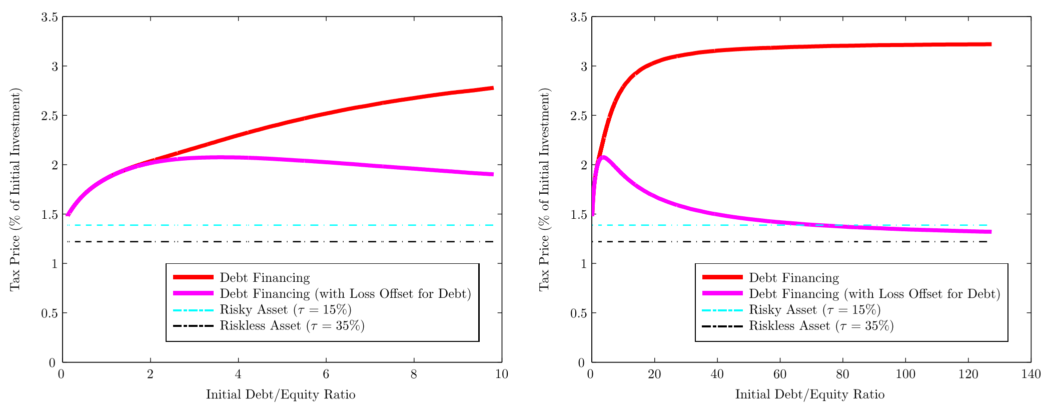

(a) Debt/Equity Ratio Less than 10 (b) Wider Range of Debt/Equity Ratios

Figure 5: The figures show the aggregate price of tax for two taxpayers investing in an asset for one year, with one taxpayer putting up the equity and the other providing debt financing. Debt is taxed at the rate of $\tau_D = 35\%$, and equity is taxed at the favorable rate of $\tau_E = 15\%$. No offset for losses is available for the equity investor, but for the debt investor, both the situation in which loss offsets are allowed and the situation in which they are not are illustrated. The underlying risky asset has a volatility of 18.9%, and the risk-free rate is 3.6%. In both graphs, the initial debt-to-equity ratio, computed on a pre-tax basis, varies along the horizontal axis. A wider range of ratios is shown in the figure on the right. The aggregate price is expressed as a percentage of aggregate initial investment value. The illustration assumes that the "debt" is not recharacterized as equity for tax purposes even when leverage ratios are very high.

It is notable that the relative positions of the burden of the T_{NL} tax on the risky asset and the riskless asset in Figure 5 are close. This is a result of the fact that the riskless asset is assumed to be taxed at the $\tau_D = 35\%$ rate for the debt owner and the risky asset is assumed to be taxed at the $\tau_E = 15\%$ rate for the equity holder. If the relative rates changed, or if the period of investment were longer than the one-year investment reflected in the figure, the relative sizes of the burdens may shift, and the asset with the higher burden could from being the risky asset to the riskless asset.

As a final comment about subadditivity and the penalty on division of ownership, I note that the T_{NL} tax would no longer be subadditive if it allowed a certain amount of loss, up to a capped amount per taxpayer. An example of such a rule is the $3,000 limit on losses for individuals under IRC § 1211. In such a case, if ownership is divided between two taxpayers and each has a loss, the aggregate limit of losses allowed for tax purposes is double what it would be in there were only one owner, and this can make division of ownership more advantageous. Figure 6 illustrates this general phenomenon and highlights the range of possible incomes for two co-owners for which divided ownership is actually preferable to exclusive ownership. The importance of this type of deviation from subadditivity depends upon the magnitude of the potential benefit relative to the size of the overall investment. If the cap on losses is small relative to the overall investment, then the version of the T_{NL} tax that allows a loss has essentially the same properties as the usual T_{NL} tax. It is possible, however, that for some taxpayers and certain situations the difference may be important. In these cases, the price of the modified T_{NL} tax can be calculated to determine the burden or benefit that comes from the tax rules when ownership is divided.

6 Extensions

It is possible to relax various assumptions underlying the model developed in Section 3. For example, it is straightforward to generalize to a setting in which consumption and tax payments may occur at multiple discrete times, ore even continuously, rather than requiring that all consumption of after-tax investment wealth occur at the end. As long as the timing and amounts of consumption, as well as the applicable tax rules, are all governed by deterministic functions of asset prices, the government's right to collect taxes over the course of the entire interval from t_0 to t_1 is simply a contingent claim that makes payments with timing and size also given by a deterministic function of asset prices. The price of such a claim can be calculated on a risk-neutral basis using replicating portfolios in the same way as a claim with only one payment date, provided that markets are sufficiently complete to allow the necessary replication. In this case, the price of the tax is once again exactly the price of the corresponding claim.

The removal of the restriction on the timing of consumption would be particularly useful in analyzing how the realization requirement for tax gains and losses impacts the price of a tax. For example, if a taxpayer followed a rule for taking losses whenever possible and putting off gains until the end of a specified time horizon, the price of the tax would generally be reduced. The extent to which deductions for losses are permitted, as well as wash sale rules and other special tax code provisions, could be factored into the analysis as well in order to develop an understanding of the true price of the tax. For my methodology to apply, it is necessary for the taxpayer's behavior to be deterministic, given a evolution of stochastic asset prices, and given the relevant tax rules. This is arguably the case, however, in many realistic situations,

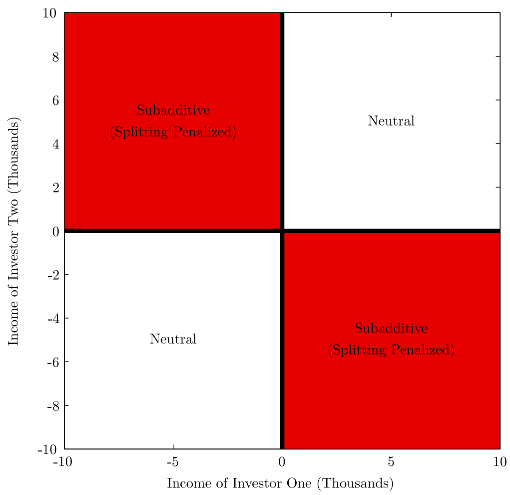

(a) No Loss Offsets

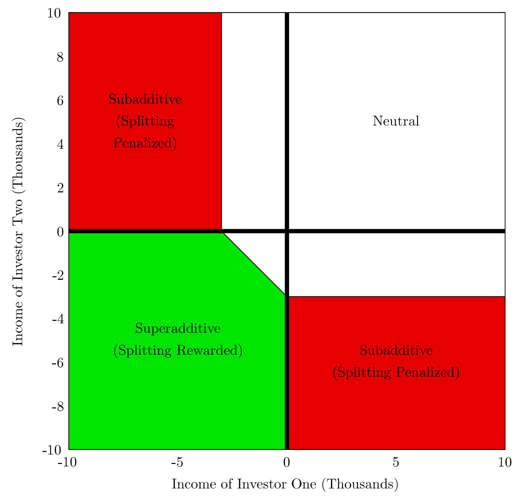

(b) Loss Offsets up to $3,000

Figure 6: The figure on the left shows the regions in which the T_{NL} tax burdens a division of ownership. The horizontal and vertical axes represent possible income outcomes for two investors who each own a portion of a single asset; together they own the entire asset. If a pair of income outcomes is in the red area, the division of ownership is more heavily burdened by the T_{NL} tax than unified ownership by a single taxpayer would be. If it is in the white area, there is no tax advantage or disadvantage to division of ownership. The figure on the right shows the similar picture for a tax that is proportionate except that it allows $3,000 of loss offsets. In this case, the red and the white areas play the same role as in the figure on the left. There is now also a green area, however, in which division of ownership is tax favored relative to unified ownership.

and I plan to pursue this line of analysis in future work applying my framework.

The model can also be generalized to incorporate a number of other possibilities, such as transaction costs, illiquid assets, and otherwise incomplete markets. Determining the price of a tax in such an extended model is still a problem that directly parallels the pricing of a government-held contingent claim, and a price can be ascertained to the extent that techniques from asset pricing theory provide a way to find a price for the claim. If perfect replication of the contingent claim with portfolios of assets in appropriate underlying securities remains possible, then an exact price for the claim, and hence the tax, is available. In many cases of incomplete markets, however, perfect replication may not be feasible. Even in these more challenging situations, it is still often possible at least to obtain upper and lower bounds for the price of the claim, and hence the price of the tax. There are various estimation techniques of this type, including, for example, the "good-deal bounds" described in Chapter 18 of Cochrane (2001), and depending upon the specific details involved, the estimates can provide a significant amount of information about the range in which the price of the claim, and the tax, must fall.

7 Conclusion

The framework I have developed extends the general equilibrium techniques of Kaplow (1991, 1994) to non-linear taxes by allowing trading and rebalancing of investment portfolios between initial investment and final consumption of wealth. Asset pricing theory applies directly and allows calculation of the price of the government's contingent tax claim with respect to a specified investment strategy for a taxpayer. This price of the claim truly represents the "price" of the tax because imposing an up-front lump-sum tax in this amount, along with appropriate dynamic investment portfolio modifications, preserves a state of general equilibrium.

My methodology provides a robust set of tools for assessing critically the nature of a tax system with non-linear imperfections. In an ideal setting, it is likely most desirable to eliminate non-linearities, but such sweeping change may not be feasible. As a second-best solution, however, some degree of incremental reform may be possible. My concept of the price of a tax provides a practical method for identifying the areas most in need of this type of reform and also sheds light on the types of reforms that will be most effective.

As an example, I applied my framework to analyze the particular case of the non-linear T_{NL} tax, which is linear for gains but does not allow any offset for losses. This example is informative because it is not much more complicated than a linear tax, but it significantly burdens assets in ways that linear taxes do not. In addition, similar taxes are actually present in the current U.S. tax system. My findings included the result that convex taxes, such as the T_{NL} tax, generally burden the risk in risky returns, but not the expected level of these returns. Accordingly, such taxes encourage portfolio diversification targeted at risk minimization without regard to expected returns, and this produces a distortion in investment allocation decisions. I also found that the T_{NL} tax imposes a particularly heavy burden on out-of-the-money call and put options, because of the high degree of risk associated with such instruments. In addition, I found that subadditive taxes, such as the T_{NL} tax, generally penalize synthetic division of risky asset ownership across taxpayers accomplished using put-call parity or debt financing.

From the perspective of guiding tax policy reform, my findings tend to give preliminary support to the ideas that it may be desirable to have lower tax rates for taxpayers with relatively riskier portfolios and also to mitigate the aggregate tax burden on assets with debt-financed or synthetically divided ownership. It is my hope that more work using the price of a tax can be undertaken to gain further insights into what reform choices would be most desirable in terms of both distributional fairness and efficiency.

A Appendix

A.1 Convexity and Subadditivity of the T_{NL} Tax

Proposition A.1. *The T_{NL} tax defined in* (7) *is subadditive and convex, as those concepts are defined in* (8) *and* (12)*, respectively.*

Proof. To see that subadditivity holds, it is easiest to consider four separate cases depending upon the signs of x and y, namely

$$T_{\mathrm{NL}}(x+y) \leq \begin{cases} 0, & \text{if } x < 0 \text{ and } y < 0; \\ \tau \max(0, y), & \text{if } x < 0 \text{ and } y > 0; \\ \tau \max(0, x), & \text{if } x > 0 \text{ and } y < 0; \text{ and} \\ \tau(x+y), & \text{if } x > 0 \text{ and } y > 0. \end{cases}$$

In each case, the expression on the right-hand side of the equation is exactly equal to the sum of $T_{\mathrm{NL}}(x)$ and $T_{\mathrm{NL}}(y)$. Thus subadditivity follows, since $T_{\mathrm{NL}}(x+y) \leq T_{\mathrm{NL}}(x) + T_{\mathrm{NL}}(y)$.

To see that convexity holds, note first that $T_{\mathrm{NL}}(\alpha x) = \alpha T_{\mathrm{NL}}(x)$ for any constant $\alpha \geq 0$. As a result, the condition for convexity given in (8) is equivalent to

$$T_{\mathrm{NL}}(\alpha x + (1-\alpha) y) \leq T_{\mathrm{NL}}(\alpha x) + T_{\mathrm{NL}}((1-\alpha) y),$$

and this inequality holds because of the subadditivity of T_{NL}. □

A.2 Proof of Proposition 1

Let A_t and B_t be the prices of the risky and riskless assets, respectively, at time t, and write $\Delta A = A_{t_1} - A_0$ and $\Delta B = B_{t_1} - B_0$ for the changes in value from the initial time $t = 0$ to the end of the investment period at time t_1. Also, write $r = \frac{1}{t_1} \log \frac{B_{t_1}}{B_0}$ for the annual risk-free rate of return. The price of the tax T for the risky asset A satisfies

$$\mathcal{P}(A; T) = e^{-rt_1} \mathrm{E}\left[T(\Delta A)\right] \geq e^{-rt_1} T\left(\mathrm{E}\left[\Delta A\right]\right) = e^{-rt_1} T(\Delta B) = \mathcal{P}(B; T), \quad \text{(A.1)}$$

where the inequality follows from Jensen's Inequality and the final equality is a result of the fact that the expectation is taken with respect to the risk-neutral distribution for returns on A, meaning that $\mathrm{E}\left[A_{t_1}\right] = B_{t_1}$.

The inequality for a concave function T is proved similarly. The fact the inequalities are strict if T is strictly convex or concave, respectively, is also a result of the usual result for Jensen's Inequality, as well as the fact that a risky asset A does not have constant returns.

A.3 Proof of Proposition 2

An approximate expression for $\mathcal{P}(A, T_{\mathrm{NL}})$ and an error term follow from the Taylor expansions for $N(x)$ and e^x. In particular, direct computation shows that

$$N(x) = \frac{1}{2} + \left(\frac{1}{\sqrt{2\pi}}\right) x + R_N(x), \quad \text{where} \quad |R_N(x)| \leq \frac{|x|^3}{2\sqrt{2\pi}},$$

and that

$$e^x = 1 + R_e(x), \quad \text{where} \quad |R_e(x)| \leq |x|, \quad \text{provided} \quad x \leq 0.$$

Since $d = \left(\frac{\sigma}{2} + \frac{r}{\sigma}\right)\sqrt{t_1}$, it thus follows that

$$\mathcal{P}(A, T_{\mathrm{NL}}) = \tau A_0 \left(N(d) - N\left(d - \sigma\sqrt{t_1}\right) e^{-rt_1}\right) = \tau A_0 \left(\frac{1}{\sqrt{2\pi}}\right)\left(\sigma\sqrt{t_1}\right) + E,$$

where

$$|E| \leq \tau A_0 \left(\frac{8}{\sqrt{2\pi}}\right)\left(\frac{\max\left(\sigma^2/2, r\right)^3}{\sigma^3}\right) t_1^{3/2} + rt_1.$$

It is straightforward to use this result to show that $|E| \leq 5\tau A_0 \sigma^2 t_1$ when $r \leq \sigma^2$ and $\sigma^2 t_1 \leq 1$, and this suffices to prove the proposition.

A.4 Proof of Proposition 3

The lower bound for $\mathcal{P}(A, T_{\mathrm{NL}})$ follows from the fact that $N(d) \geq 1 - N\left(d - \sigma\sqrt{t_1}\right)$, since $d > 0$ and $\sigma\sqrt{t_1} \leq 2d$. Plugging this inequality for N into (9) shows that

$$\mathcal{P}(A, T_{\mathrm{NL}}) \geq \tau A_0 \left(1 - N\left(d - \sigma\sqrt{t_1}\right)\left(1 - e^{-rt_1}\right)\right),$$

and this is the desired lower bound. The upper bound for $\mathcal{P}(A, T_{\mathrm{NL}})$ follows directly from the fact that $N(x) \leq 1$ for all x.

In the case of the riskless asset, the formula in (9) is no longer applicable, since it requires a positive volatility. Instead, the price of the linear tax from (5) can be used, since T_{NL} and T_{L} are the same for the riskless asset. This formula shows directly that

$$\mathcal{P}(B, T_{\mathrm{NL}}) = \tau A_0 \left(1 - e^{-rt_1}\right),$$

and the desired lower and upper bounds both follow immediately from this equality.

A.5 Proof of Proposition 4

It follows directly from the definition of the price of a tax in (3) and the definitions of T_{L} and T_{NL} in (4) and (7) that

$$\mathcal{P}\left(A, T_{\mathrm{NL}}\right) = \mathcal{P}\left(A, T_{\mathrm{L}}\right) + \tau e^{-rt_1} \mathrm{E}\left[\max\left(0, A_0 - A_{t_1}\right)\right].$$

If there is a positive probability that $A_{t_1} < A_0$, then the second term on the right-hand side of this equation is positive, and otherwise it is zero. This suffices to prove the proposition.

A.6 Proof of Proposition 5

Let r be the risk-free rate of return, let $C_t(K)$ be the price of a call option with strike K at time t, and write W for the initial amount of wealth invested in call options with strike K. The tax burden on this investment is

$$\begin{aligned}
\mathcal{P}\left(C(K), T_{\mathrm{NL}}\right) &= e^{-r}\mathrm{E}\left[T_{\mathrm{NL}}\left(W\left(\frac{C_1(K) - C_0(K)}{C_0(K)}\right)\right)\right] \qquad \text{(A.2)}\\
&= \tau W e^{-r}\mathrm{E}\left[\max\left(0, \frac{C_1(K) - C_0(K)}{C_0(K)}\right)\right]\\
&= \tau W e^{-r}\mathrm{E}\left[\max\left(0, \frac{\max\left(0, A_1 - K\right) - C_0(K)}{C_0(K)}\right)\right],\\
&= \tau W e^{-r}\mathrm{E}\left[\max\left(0, \frac{A_1 - K - C_0(K)}{C_0(K)}\right)\right],
\end{aligned}$$

where the second line follows from the definition of T_{NL}, the third line follows from the definition of $C_1(K)$, and the final line follows from the fact that the argument of the expectation operator in the third line is only positive in the range $A_1 > K + C_0(K)$.

It is convenient to introduce notation for some relevant intervals, namely:

$$R_1 = \left[K, K + C_0(K)\right], \quad R_2 = \left[K + C_0(K), \infty\right), \quad \text{and} \quad R_3 = R_1 \cup R_2 = \left[K, \infty\right).$$

Also, let $\mathbf{1}_R$ be the characteristic function of an interval R, i.e., the function that has the value 1 inside the interval and zero elsewhere. The derivative with respect to strike price of the tax burden on the call option investment is

$$\begin{aligned}\frac{\partial \mathcal{P}\left(C(K), T_{\mathrm{NL}}\right)}{\partial K} &= -\tau W e^{-r}\mathrm{E}\left[\mathbf{1}_{R_2}\left(\frac{1}{C_0(K)} + \frac{A_1 - K}{C_0^2(K)/C_0'(K)}\right)\right] \\ &= -\tau W e^{-r}\left(\frac{-e^r C_0'(K) - \mathrm{E}\left[\mathbf{1}_{R_1}\right]}{C_0(K)} + \frac{e^r C_0(K) - \mathrm{E}\left[\mathbf{1}_{R_1}(A_1 - K)\right]}{C_0^2(K)/C_0'(K)}\right) \\ &= \tau W e^{-r}\left(\frac{C_0(K)\mathrm{E}\left[\mathbf{1}_{R_1}\right] + C_0'(K)\mathrm{E}\left[\mathbf{1}_{R_1}(A_1 - K)\right]}{C_0^2(K)}\right) \\ &\geq \tau W e^{-r}\left(\frac{\mathrm{E}\left[\mathbf{1}_{R_1}\right](1 + C_0'(K))}{C_0(K)}\right) \geq 0.\end{aligned} \tag{A.3}$$

The first line of (A.3) follows from direct computation of the derivative. The second line of (A.3) makes use of the identities

$$C_0(K) = e^{-r}\mathrm{E}\left[\mathbf{1}_{R_3}(A_1 - K)\right] \quad \text{and} \quad C_0'(K) = -e^{-r}\mathrm{E}\left[\mathbf{1}_{R_3}\right]. \tag{A.4}$$

The third line of (A.3) is simply an algebraic rearrangement of the second line. The fourth line of (A.3) makes use of the inequalities

$$\mathrm{E}\left[\mathbf{1}_{R_1}(A_1 - K)\right] \leq C_0(K)\mathrm{E}\left[\mathbf{1}_{R_1}(A_1 - K)\right] \quad \text{and} \quad |C_0'(K)| = e^{-r}\mathrm{E}\left[\mathbf{1}_{R_3}\right] \leq 1,$$

with the first of these resulting from the fact that $A_1 - K \leq C_0(K)$ throughout the interval R_1 and the second of these resulting from the fact that $r \geq 0$. As a result of (A.3), it follows that the derivative is always non-negative, and so the tax burden is increasing with K, as desired.

To find the limit of the tax burden on the call investment as $K \to \infty$, note that

$$\begin{aligned}\mathrm{E}\left[\mathbf{1}_{R_3}\left(\frac{A_1 - K}{C_0(K)}\right)\right] - \mathrm{E}\left[\mathbf{1}_{R_2}\left(\frac{A_1 - K - C_0(K)}{C_0(K)}\right)\right] &= \mathrm{E}\left[\mathbf{1}_{R_1}\left(\frac{A_1 - K}{C_0(K)}\right)\right] + \mathrm{E}\left[\mathbf{1}_{R_2}\right] \\ &\leq \mathrm{E}\left[\mathbf{1}_{R_1}\right] + \mathrm{E}\left[\mathbf{1}_{R_2}\right] \to 0,\end{aligned}$$

where the final limit limit is taken as $K \to \infty$. This result and the expression for the tax burden in the last line of (A.2) combine to show that

$$\lim_{K\to\infty} \mathcal{P}\left(C(K), T_{\mathrm{NL}}\right) = \tau W e^{-r}\mathrm{E}\left[\mathbf{1}_{R_3}\left(\frac{A_1 - K}{C_0(K)}\right)\right] = \tau W,$$

where the last equality follows from the expression for $C_0(K)$ in (A.4). This is the result claimed in the statement of the proposition.

At the other extreme, as $K \to 0^+$, $C_0(K) \to A_0$, and so the final expression for

$\mathcal{P}\left(C(K), T_{\mathrm{NL}}\right)$ in (A.2) shows that

$$\lim_{K\to 0^+} \mathcal{P}\left(C(K), T_{\mathrm{NL}}\right) = \tau W e^{-r} \mathrm{E}\left[\max\left(0, \frac{A_1 - A_0}{A_0}\right)\right] = \mathcal{P}\left(A, T_{\mathrm{NL}}\right),$$

and this is simply the tax burden on an investment of an amount W in the risky asset, as claimed in the proposition.

The results stated in the proposition for puts may be proven in a similar fashion to those for calls.

A.7 Proofs of Proposition 6 and Corollary 1

Proof of Proposition 6. Let D_t and E_t be the prices of the debt and equity positions at time t, and let $A_t = D_t + E_t$ be the price of the asset at time t. Also, write

$$\Delta D = D_1 - D_0, \quad \Delta E = E_1 - E_0, \quad \text{and} \quad \Delta A = A_1 - A_0$$

for the changes in value of each quantity from time $t = 0$ to time $t = 1$, and write r for the risk-free rate of return. If T is subadditive, it follows that

$$\begin{aligned} \mathcal{P}(T; D) + \mathcal{P}(T; E) &= e^{-r_B} \mathrm{E}\left[T\left(\Delta D\right)\right] + e^{-r_B} \mathrm{E}\left[T\left(\Delta E\right)\right] \\ &= e^{-r_B} \mathrm{E}\left[T\left(\Delta D\right) + T\left(\Delta E\right)\right] \\ &\geq e^{-r_B} \mathrm{E}\left[T\left(\Delta D + \Delta E\right)\right] \\ &= e^{-r_B} \mathrm{E}\left[T\left(\Delta A\right)\right] = \mathcal{P}(T; A), \end{aligned} \tag{A.5}$$

where the second line follows from the linearity of the expectation, the third line follows from the subadditivity of T, and the fourth line follows from the definitions of A, D and E.

If T is superadditive, the opposite inequality can be proven in a similar fashion. □

Proof of Corollary 1. To obtain the result, it is necessary to analyze the behavior of the signs of the quantities ΔD and ΔE introduced in the proof of Proposition 6 With respect to debt, note that the fact that the distribution for A_1 has support in the region $K < A_1$ implies that the value of D_0 must be less than K, since it is simply the present value of the expectation of $\min(K, A_1)$. As a result, ΔD, when viewed as a function of A_1, monotonically increases and switches sign from negative to positive at some value $A_1^D < K$. With respect to equity, the value of E_0 must be positive, since it is the present value of $\max(0, A_1 - K)$ and A_1 has support in the region $A_1 > K$. As a result, the function ΔE, when viewed as a function of A_1, is also monotonically increasing and switches sign from negative to positive at some value $A_1^E > K$.

From the foregoing results, it is clear that ΔD is positive and ΔE is negative in the interval $\left[A_1^D, A_1^E\right]$, and that $A_1^D < K < A_1^E$. Because A_1 has support in a neighborhood of K, it thus follows that A_1 has support in an interval in which ΔD and ΔE have opposite signs. Throughout this interval, there is a strict inequality $T_{\mathrm{NL}}(\Delta D) + T_{\mathrm{NL}}(\Delta E) > T_{\mathrm{NL}}(\Delta D + \Delta E)$, and so there is strict inequality of expectations in (A.5) in the proof of Proposition 6. This proves the corollary. □

References

Atkinson, A. B. and Stiglitz, J. E., 1980, *Lectures on Public Economics*, McGroaw-Hill, London.

Avi-Yonah, R. S., 2004, "Risk, Rents and Regressivity: Why the United States Needs Both an Income Tax and a VAT," *Tax Notes*, 105, 1651.

Black, F. and Scholes, M., 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 3, 637–654.

Bodie, Z., Kane, A., and Marcus, A. J., 2005, *Investments*, McGraw-Hill Irwin, New York, 6th edition.

Campisano, M. and Romano, R., 1981, "Recouping Losses: The Case for Full Loss Offsets," *Northwestern University Law Review*, 76(5), 709–744.

Cochrane, J. H., 2001, *Asset Pricing*, Princeton University Press, Princeton, New Jersey.

Cox, J. C., Ross, S. A., and Rubinstein, M., 1979, "Option Pricing: A Simplified Approach," *Journal of Financial Economics*, 7, 229–263.

Cunningham, N. B., 1996, "The Taxation of Capital and the Choice of Tax Base," *Tax Law Review*, 52, 17–44.

Domar, E. D. and Musgrave, R. A., 1944, "Proportional Income Taxation and Risk-Taking," *Quarterly Journal of Economics*, 58, 388–422.

Domar, E. D. and Musgrave, R. A., 1945, "Income Taxation and Risk Taking," *TAXES*, 23, 60–62.

Gentry, W. M. and Hubbard, R. G., 2005, ""Success Taxes," Entrepreneurial Entry, and Innovation," in A. B. Jaffe, J. Lerner, and S. Stern, (eds.) "Innovation Policy and the Economy," volume 5, pp. 87–108, MIT Press.

Green, R. C. and Talmor, E., 1985, "The Structure and Incentive Effects of Corporate Tax Liabilities," *The Journal of Finance*, 40, 1095–1114.

Harrison, J. M. and Kreps, D. M., 1979, "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory*, 20, 381–408.

Kaplow, L., 1991, "Taxation and Risk Taking: A General Equilibrium Perspective," *NBER Working Paper*, 3709.

Kaplow, L., 1994, "Taxation and Risk Taking: A General Equilibrium Perspective," *National Tax Journal*, 47, 789–798.

Majd, S. and Myers, S. C., 1986, "Tax Asymmetries and Corporate Income Tax Reform," *NBER Working Paper*, 1924.

Merton, R. C., 1973, "The Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4, 141–183.

Merton, R. C., 1974, "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *The Journal of Finance*, 29(2), 449–470.

Merton, R. C., 1976, "Option Pricing When Underlying Stock Returns are Discontinuous," *Journal of Financial Economics*, 3, 125–144.

Sandmo, A., 1985, "The Effects of Taxation on Savings and Risk Taking," in A. J. Auerbach and M. S. Feldstein, (eds.) "Handbook of Public Economics," pp. 265–311, North Holland, New York.

Schenk, D. H., 2000, "Saving the Income Tax With a Wealth Tax," *Tax Law Review*, 53, 423–475.

Stiglitz, J. E., 1969, "The Effects of Income, Wealth, and Capital Gains Taxation on Risk-Taking," *Quarterly Journal of Economics*, 83, 263–283.

Stiglitz, J. E., 1972, "Taxation, Risk Taking, and the Allocation of Investment in a Competitive Economy," in M. C. Jensen, (ed.) "Studies in the Theory of Capital Markets," Praeger, New York.

Warren, A. C., 1996, "How Much Capital Income Taxed Under an Income Tax Is Exempt Under a Cash Flow Tax?" *Tax Law Review*, 52, 1–16.

Weisbach, D. A., 2004a, "The (Non)Taxation of Risk," *Tax Law Review*, 58, 1–57.

Weisbach, D. A., 2004b, "Taxation and Risk-Taking with Multiple Tax Rates," *National Tax Journal*, 57(2), 229–243.

Zelenak, L., 2006, "The Sometimes-Taxation of the Returns to Risk-Bearing Under a Progressive Income Tax," *SMU Law Review*, 59, 879–915.