

**OPTIMAL TAX PORTFOLIOS
AN ESTIMATION OF GOVERNMENT TAX REVENUE
MINIMUM-VARIANCE FRONTIERS**

NATHAN SEEGER

ABSTRACT. Government budget crises in the 2000s were magnified by the increase in tax revenue volatility governments experienced. Governments can decrease the variance of their tax revenues by holding the efficient “portfolio” of taxes. In this conceptualization, each tax base is a potential asset the government can hold and the tax rate on a given base is the weight they put on the asset. Conceptualizing government finances as an optimal portfolio problem highlights the ability of governments to hedge risk by taxing different bases, but this method must be adapted to account for numerous differences between a government and an individual investor. This paper conducts the analysis of the mean-variance tradeoff made by governments within a utility framework. This analysis demonstrates the tradeoffs governments face between volatility and deadweight loss and between public and private consumption volatility. Therefore, the government does not minimize tax revenue volatility but aims to optimize tax revenue volatility. As an application of the theoretical model, I create a method to estimate the minimum variance a government can achieve for a given expected level of tax revenue: an efficient frontier. I demonstrate the method with a few examples using data from U.S. state governments. Estimating state-specific efficient frontiers allows for across state analysis of the relative mean-variance tradeoffs. In addition, the different portfolios states have held can be plotted to determine how state portfolios have changed over time relative to the efficient frontier.

JEL Numbers: H21, H7, H68, R51

I would like to thank seminar participants at the U.S. Treasury and the Michigan Public Finance Seminar for providing valuable feedback, and in particular Jim Hines, David Albouy, Steve Salant, Scott Page, Joel Slemrod, and Dan Silverman for their help and advice. Any mistakes are my own. Please contact the author at seegert@umich.edu.

Government budget crises in the 2000s were magnified by an increase in tax revenue volatility. For example, state governments in the United States experienced a 500 percent increase in volatility in the 2000s relative to previous decades. State governments are particularly sensitive to tax revenue volatility because of their balanced budget rules and other frictions which make smoothing tax revenues difficult. Governments can decrease the variance of their tax revenues by holding the efficient “portfolio” of taxes. In this conceptualization, each tax base is a potential asset held by the government and the tax rate on a given base is the weight the government puts on the asset. Conceptualizing government finances as an optimal portfolio problem highlights the ability of governments to hedge risk by taxing different bases, but this method must be adapted to account for numerous differences between a government and an individual investor.

The first of these differences results from the obvious disparity in size between individual investors and governments. When an individual investor increases her holdings of a given asset, the mean return of the asset is not affected. However, when the government increases the weight on a given asset by increasing the tax rate, the asset’s mean return decreases because the tax base shrinks as a result of individual behavioral responses to the tax increase. The decreased return is the leakage caused by behavioral responses by individuals.

The second difference occurs because, in contrast to the individual investor, the assets in the government’s portfolio are interdependent. When a government increases its income tax rate, this affects sales tax and corporate tax returns. For example, individual behavioral responses of how much to consume and how much income to shift between income and corporate tax bases depends on the income tax rate. The effect of a given tax rate on other tax bases is a horizontal externality that complicates government finance.

The third difference occurs because, in contrast to the individual investor, the government’s objective function is not to minimize the variance of tax revenue for a given expected rate of return but to maximize expected utility. One aspect that needs to be considered when maximizing

expected utility is the cost from volatile tax revenue streams. Other aspects the government must consider include costs due to deadweight loss and the efficient risk-sharing between public and private consumption.

This paper conducts the analysis of the mean-variance tradeoff made by governments within a utility framework. The analysis demonstrates the tradeoffs governments face between volatility and deadweight loss and between public and private consumption volatility. Therefore, the government aims to optimize, not minimize, tax revenue volatility.

The theoretical model is applicable to governments that are constrained, in some way, from perfectly smoothing their revenue causing their expenditures to be exposed to risk in revenue. For example, U.S. state governments are limited in their ability to smooth revenue because of self-imposed balanced budget rules. In addition, European governments' expenditures are currently exposed to additional risk in revenues because of limits in borrowing caused by the Euro-zone debt crises. Borrowing constraints can also cause developing countries' expenditures to be exposed to risk in their revenue streams. In fact, all governments are exposed to revenue shocks to some extent because of their uncertainty about whether an observed shock is temporary or permanent. Therefore, explaining tax revenue volatility and the ways in which tax policy can be used to stabilize government expenditures is important for governments world-wide.

As an application of the theoretical model, I create a method to estimate the minimum variance governments can achieve for a given expected level of tax revenue: a minimum-variance frontier. Following the results of the theoretical model, each tax base-rate pair is considered a separate asset. To implement this method, counterfactual portfolio returns first had to be estimated because data exist for only one portfolio in any given year (the actual portfolio held by the government).

I demonstrate the method with a few examples using data from U.S. state governments. Estimating state-specific minimum-variance frontiers allows for across-state analysis of the relative

mean-variance tradeoffs. In addition, the different historic portfolios held by governments can be plotted to determine how government portfolios have changed over time relative to the minimum-variance frontier.

1. LITERATURE

This paper builds upon a large literature of optimal taxation, which has progressed by isolating new tradeoffs involved with taxation. Mossin (1968b) and Stiglitz (1969) study the effect of taxes on risk taking by individuals.¹ Mirrlees' 1971 seminal paper presented the tradeoff between equity and efficiency inherent in progressive income taxation. Allingham and Sandmo (1971) discuss the problem of tax evasion, which they describe as a new tradeoff that is "of considerable practical interest." The potential for taxation to act as social-insurance was demonstrated by Varian in his 1980 paper.² Stern's 1992 paper considers new tradeoffs that exist in a dynamic setting and Slemrod's 1990 paper discusses the new tradeoffs of taxation when the entire system of taxation is considered.

This paper is also influenced by the tax portfolio literature started by Groves and Kahn in 1952. This literature has focused on the income elasticities and stability of state and local taxes.³ Dye and McGuire (1991) compare state sales and income taxes in an attempt to determine which tax base is more stable, but find that their stabilities cannot be systematically differentiated. Sobel and Holcombe (1996) extend Dye and McGuire's analysis by including more tax instruments in a new time series technique but find similar ambiguities. Bruce, Fox, and Tuttle (2006) use disaggregated data to refine the literature's results; for example, they use actual tax base data, rather than proxies, to estimate the relationship between personal income and tax bases.

¹Mossin (1968b) and Stiglitz (1969) study the effect of taxes on risk taking by individuals. In contrast, here I consider the effect of taxes on the exposure of risk to state tax revenues.

²Varian (1980) describes the insurance benefit of taxation that occurs when luck can cause differences in income; in contrast, this paper is focused on the portfolio choice of governments that take into account the costs of exposing themselves to more risk.

³This literature has noted that the elasticities change over time (Groves and Kahn (1952)) and with the business cycle (Fox and Campbell (1984) and Otsuka and Braun (1999)).

Their results suggest that neither the personal income tax nor the sales tax is universally more volatile than the other. Instead of comparing individual tax bases, my paper uses adapted optimal portfolio theory to demonstrate that a mix of tax bases may provide the most stable tax revenues.

2. MODEL SETUP

A *Timing* The economy is assumed to be a one period snapshot of a dynamic model where the state of nature within the period is uncertain *ex ante*. The timing is given below, but note the government moves before the state of nature is realized, causing uncertainty from the government's point of view; in contrast, the individual does not face uncertainty because she moves after the state of nature is realized. The government does not know the realization of the wage and profit, but is assumed to know the distribution.

<i>Order of Decisions</i>	<i>Choices</i>
1st - Government	τ_c, τ_w
2nd - Nature	w, π
3rd - Individual	c, L, β
4th - Production occurs	
5th - Utility realized	

B. *Individual Behavior.* The individual has utility over her supply of labor L , the public good g , and total private consumption c , which is split between goods that are taxed, $c_1 \equiv \beta c$, and goods that are untaxed $c_2 \equiv (1 - \beta)c$. The individual chooses c , L , and β to maximize utility

$$\max_{c, \beta, L} \quad u = U(c, \beta, L, G,) \quad \text{subject to} \quad y = c(1 + t_c \beta)$$

where t_c is the tax rate on consumption and $y = (1 - t_w)wL + \pi$ is income net of the wage income tax.⁴ Wages and profits are assumed to be stochastic, resulting in stochastic consumption and wage income. Consumption and its mean and variance can be written as,

$$c = (1 - \tau_c \beta)((1 - t_w)wL + \pi) \quad \text{where} \quad \tau_c = t_c / (1 - t_w \beta)$$

$$\bar{c} = (1 - \tau_c \beta)((1 - t_w)\bar{w}\bar{L} + \bar{\pi}) \quad \sigma_c^2 = (1 - \tau_c \beta)^2((1 - t_w)^2 \sigma_{wL}^2 + \sigma_\pi^2 + 2(1 - t_w)\sigma_{wL,\pi})$$

Consumption and wage income will not be perfectly correlated as long as wages and profits are not perfectly correlated, which can be seen in figure 1. Figure 1 represents consumption as a vector equal to the sum of the vectors of wage and profit income where the lengths of all of the vectors equal the standard deviation of the variable. Using the law of cosines, the correlation between two vectors is depicted as the cosine of the angle between any two vectors. For example, if the vectors are parallel the variables are perfectly correlated and if the vectors are perpendicular the variables are independent.

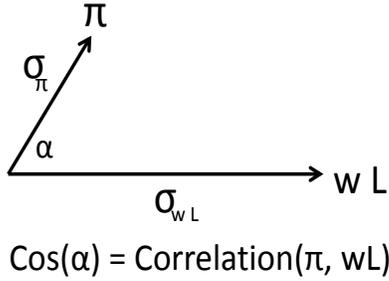
The ability of the government to hedge idiosyncratic risk between consumption and wage income tax bases depends on the correlation of these two variables.⁵ In this example if the standard deviation of the profit shock increases, the correlation between consumption and wage income decreases. This can be seen graphically by increasing the length of the profit vector extending from the end of the wage income vector, which results in a larger angle between consumption and wage income (and also decreases the cosine of the angle and therefore the correlation).⁶

⁴ $wL(1 - t_w) + \pi = y = (1 - \beta)c + \beta c(1 + t_c) = c(1 + t_w \beta)$

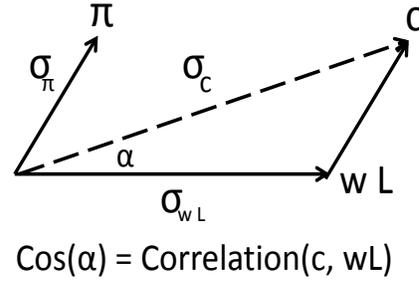
⁵First, let $\tau_c = t_w = 0$ for simplicity, allowing $c = wL + \pi$. The cosine of the angle between wage income and consumption, using the law of cosines, can be written as $\cos(\theta) = (\sigma_c^2 + \sigma_{wL}^2 - \sigma_\pi^2) / (2\sigma_{wL}\sigma_c)$. The numerator can be reduced to $2\text{cov}(wL, c)$ using the variance formula $\text{var}(\pi) = \text{var}(c - wL) = \text{var}(c) + \text{var}(wL) - 2\text{cov}(wL, c)$. Therefore the cosine of the angle between wage and profit income is equal to the correlation between them; $\cos(\theta) = \text{cov}(wL, c) / (\sigma_{wL}\sigma_c) = \rho_{wL,c}$.

⁶In this example increasing the standard deviation of profit income increases the standard deviation of consumption. However, if profit income and wage income were negatively correlated, increasing the standard deviation of profit income could decrease the standard deviation of consumption.

FIGURE 1. Vector Representation of Shocks



(a) Income and Savings Shocks



(b) Correlation Consumption and Income

Utility maximization requires that: i) the marginal disutility from supplying labor equals the marginal utility of the income it produces and ii) the ratio of marginal utilities from total consumption c and β is equal to the consumption tax rate multiplied by net income. When the consumption tax rate is zero there is no distortion between consumption goods, and the marginal utility with respect to β is zero. Composing utility in terms of total consumption, c , and β simplifies the composition of deadweight loss because β encompasses all behavioral responses between goods.

$$(1) \quad U_1(c, \beta, L, G)(1 - \tau_c\beta)(1 - t_w)w = U_3$$

$$(2) \quad \frac{U_2}{U_1} = \tau_c((1 - t_w)wL + \pi)$$

C. *Government* The government produces the public good g and finances its production with taxes on consumption and wage income. Both of the government's tax bases are state-dependent and uncertain when the government makes its decision. Therefore, the government maximizes the expected utility of the representative individual, which generally can be written as a function

of the moments (e.g. mean, variance, skewness) of the state-dependent variables. This analysis restricts attention to the cases where expected utility is fully characterized by a function of the first two moments (mean and variance), but is robust to considering higher moments. This restriction holds if the joint distribution of the state-dependent variables is normal or any distribution fully characterized by the first two moments (e.g. log-normal and uniform).⁷ The utility function is assumed to be additively separable such that $U_{1,4} = 0$ so the level of social welfare can be written as

$$(3) \quad E[u] = \int U(c, g) f(c, R, \sigma_R^2, \sigma_c^2) \equiv M(C, \sigma_c^2, \beta, L) + G(R, \sigma_R^2)$$

$$M_1 \geq 0, G_1 \geq 0, M_4 \leq 0, M_2 \leq 0, G_2 \leq 0$$

$$\sigma_g^2 = t_c^2 \beta^2 \sigma_c^2 + t_w^2 L^2 \sigma_w + 2(t_c \beta t_w t_w L \sigma_{w,c})$$

where \bar{c} and \bar{R} are the mean levels of the private and public consumption, σ_c^2 and σ_R^2 are the variances of private and public consumption respectively, and $G(\cdot)$ represents the expected utility from public consumption and $M(\cdot)$ represents the expected utility from private consumption. The shape of $M(\cdot)$ can differ from the shape of $G(\cdot)$, allowing for different attitudes towards volatility in public and private consumption. Specifically, the relative magnitudes of M_2 and G_2 determine the relative benefit of stable public or private consumption. Therefore, even though the government and individual are not allowed to smooth shocks through saving, the relative magnitudes of M_2 and G_2 can be thought of as the relative ability of the government and individual to smooth public and private consumption. For example, if the individual is able to perfectly smooth private consumption her utility could be written as being linear in private consumption and $M_2 = 0$.

⁷Assuming the representative individual's utility function is quadratic is another example of when the expected utility function would be characterized fully by the first two moments of the state dependent variables and is used frequently in the finance literature.

3. TAX PORTFOLIO ANALYSIS

The government's optimal portfolio problem differs from traditional optimal portfolio analysis in three important ways. First, the government is a large relative to the market, meaning the weight it puts on an asset affects the asset's returns along with other assets' returns. For example, as the government increases its tax rate on wage income, individuals supply less labor causing the wage income tax base to decrease. In addition, as the government increases its tax rate on wage income, the individual has less income to consume thus decreasing the consumption tax base. Second, the government, in maximizing the representative individual's utility, considers the welfare cost of volatility in both public and private consumption. Finally, the government, in maximizing the representative individual's utility, trades off the costs from volatility (in both public and private consumption) with deadweight loss caused by the government using taxes that distort individual's behavior.

This section begins with the general government's optimal portfolio problem which is written as minimizing the welfare cost of public and private volatility. The analysis then turns to a series of special cases that demonstrate the additional complexities in the government's optimal portfolio problem. In all of the special cases the government is constrained to bring in a given level of mean revenue, abstracting away from differences in first moments (deadweight loss). First, the government's optimal portfolio problem is considered when the representative individual is risk-neutral with respect to the private good. In this case the government's objective function reduces to minimizing the variance of the public good.

In the second case the government's optimal portfolio problem is considered when the representative individual has the same risk attitude over public and private consumption. In this case the government's objective function reduces to minimizing the sum of the variances of public and private consumption. Finally, the government's optimal portfolio problem is considered when the representative individual's risk attitude differs between public and private consumption. This

case differs from the full problem only by abstracting from changes in the first moments of public and private consumption.

Full Government's Optimal Portfolio Problem In the full model the government minimizes the negative welfare impact of volatility with endogenous levels of public and private consumption. In other words, the government minimizes the negative of utility which is the dual of the Ramsey problem from Seegert (2012a). Equation (4) provides the condition for the optimal pair of consumption and wage income tax rates. The condition demonstrates the three key differences of the government's optimal portfolio analysis from traditional optimal portfolio analysis.

First, the elasticities of the variance of public consumption with respect to a tax rate encompasses the ways in which the tax bases change according to the government's weight on each of the tax bases. Second, the welfare cost of both public and private consumption volatility is captured by the numerators of the two terms on the left-hand side of equation (4). The welfare weights ω_M and ω_G weight the elasticities in the numerator based on the risk preferences between public and private consumption. For example, if the representative individual prefers stable private consumption over stable public consumption $\omega_M > \omega_G$ and if the government is sufficiently better at smoothing tax revenues than the representative individual is at smoothing her private consumption $\omega_M < \omega_G$. Finally, the tradeoff with deadweight loss is demonstrated by the two terms on the right-hand side.

$$(4) \quad \begin{aligned} & \min_{t_w, \tau_c} - [M(c, \sigma_c^2) + G(R, \sigma_R^2)] \\ & \frac{\omega_M \varepsilon_{\sigma_c^2, \tau_c} + \omega_G \varepsilon_{\sigma_R^2, \tau_c}}{\varepsilon_{R, \tau_c}} - \frac{\omega_M \varepsilon_{\sigma_c^2, t_w} + \omega_G \varepsilon_{\sigma_R^2, t_w}}{\varepsilon_{R, t_w}} = \underbrace{\frac{\omega_\beta \varepsilon_{\beta, \tau_c} + \omega_L \varepsilon_{L, \tau_c}}{\varepsilon_{R, \tau_c}} - \frac{\omega_\beta \varepsilon_{\beta, t_w} + \omega_L \varepsilon_{L, t_w}}{\varepsilon_{R, t_w}}}_{\text{Deadweight Loss}} \end{aligned}$$

The condition in equation (4) can be broken up into four parts; two depicting the marginal costs of the consumption tax and two the wage income tax (the first and second terms of both sides

respectively) and two the costs from volatility and two the costs from deadweight loss (the left-hand side and the right-hand side respectively). The two parts on the left-hand side quantify the marginal costs due to changes in the variance of public and private consumption. The first term on the left hand side is made up of three elasticities. The elasticities in the numerator are the elasticity of the variances of public and private consumption with respect to the consumption tax rate. These elasticities are weighted by their marginal importance on utility. For example, $\omega_M = -M_2(\sigma_c^2/R)$ where $-M_2 = \partial M(C, \sigma_c^2, \beta, L)/\partial \sigma_c^2$ is positive and equal to the marginal welfare of changes in the variance of private consumption. The elasticity in the denominator is the elasticity of tax revenue with respect to the consumption tax rate. Therefore together this first term provides the marginal welfare cost of increasing the consumption tax rate on the variance of public and private consumption relative to its change in tax revenue. Similarly, the second term provides the marginal welfare cost of increasing the wage income tax rate on the variance of public and private consumption relative to its change in tax revenue.

The right-hand side of equation (4) is the welfare cost due to deadweight loss of the consumption and wage income tax rates. The first term on the right-hand side is the marginal change in deadweight loss of increasing the consumption tax rate relative to the change in tax revenue. Intuitively, increasing the consumption tax rate affects the individual's decision on which goods to consume (β) and how much labor to supply (L). The welfare cost of these changes are quantified by the numerator of the first term on the right-hand side where $\omega_\beta = -M_3(\beta/R)$ is the welfare weight applied to the elasticity of β with respect to the consumption tax rate and $\omega_L = -M_4(L/R)$ is the welfare weight applied to the elasticity of labor with respect to the consumption tax rate. Similarly, the second term on the right-hand side is the marginal change in deadweight loss of increasing the wage income tax rate relative to its change in tax revenue.

The government's efficient pair of consumption and wage income tax rates equalize the marginal welfare costs due to the variance of public and private consumption (the left-hand side

of equation (4)) with the marginal welfare costs due to deadweight loss (the right-hand side of equation (4)).

C.1 Risk Neutral Preferences with respect to Private Consumption. In this case the government's objective function reduces to minimizing the variance of tax revenue. Here, the government is constrained to producing a mean level of revenue which abstracts from first moment considerations.

$$(5) \quad \min_{t_w, \tau_c} \sigma_R^2 = t_c^2 \beta^2 \sigma_c^2 + t_w^2 L^2 \sigma_w^2 + 2t_w L t_c \beta \sigma_{c,w}$$

$$\text{Subject to: } \bar{R} = \tau_c \beta \bar{c} + t_w L \bar{w}$$

This optimization differs from typical optimal portfolio analysis because an individual investor is a small player while the government is a large player. When an individual investor increases her holdings of a given asset the mean return of the asset is not affected. When the government increases its weight on a given asset, by increasing the tax rate, the asset's mean return decreases. The decrease is the leakage caused by behavioral responses by individuals. Similarly, because the government is a large player, the weight the government places on a given asset affects the other assets in its portfolio as well. For example, lowering the income tax rate may induce individuals to shift income from the corporate tax base to the income tax base. Therefore the government must consider the ways asset returns change due to a change in the tax rate. For this example, some of these terms are assumed to simplify. Specifically, profit, wages, and labor are assumed to be independent of the consumption tax rate and β is assumed to be independent of the wage income tax rate.

The first-order conditions given in equations (6) and (7) demonstrate the additional complexity of the government's optimal portfolio analysis. In traditional portfolio analysis the assets' returns do not change with the weights, and hence only the direct effect would be included. The government's optimal portfolio problem must consider the ways asset returns change with the

weight due to leakage, horizontal externalities, and second moment effects. In equation (6) the leakage is captured by including the elasticity of labor supply with respect to the wage income tax rate. The second term represents the horizontal externality, where λ is the Lagrangian multiplier. This term considers the change in consumption due to a change in the wage income tax rate. Finally, the ways the variances and covariance of consumption and wage income change, due to a change in the wage income tax rate, is captured by the second moment effects.

$$(6) \quad \frac{\partial \sigma_R^2}{\partial t_w} = \underbrace{(2t_w L^2 \sigma_w^2 + 2t_c \beta \sigma_{w,c} - \lambda w L)}_{\text{Direct}} (1 + \underbrace{\varepsilon_{L,t_w}}_{\text{Leakage}}) + \underbrace{-\lambda \frac{t_c}{t_w} \beta c \varepsilon_{c,t_w}}_{\text{Horizontal}} + \underbrace{t_c^2 \beta^2 \frac{\partial \sigma_c^2}{\partial t_w} + t_w^2 L^2 \frac{\partial \sigma_w^2}{\partial t_w} + 2t_c \beta t_w L \frac{\partial \sigma_{w,c}}{\partial t_w}}_{\text{2nd moment}}$$

$$(7) \quad \frac{\partial \sigma_R^2}{\partial t_c} = (2t_c \beta^2 \sigma_c^2 + 2\beta t_w L \sigma_{w,c} - \lambda \beta c)(1 + \varepsilon_{\beta,t_c})$$

The optimal wage income and consumption tax rates can be found from these first-order conditions. If the government does not consider the leakage, horizontal externalities, and second moment effects associated with each tax base, the government will incorrectly set the tax rates given by t_w^* and t_c^* . However, if the government does consider the leakage, horizontal externalities, and second moment effects, the government will optimally set the tax rates given by t_w^{**} and t_c^{**} where $\theta_w = 1 + \varepsilon_{w,t_w}$ which is one plus the elasticity of the wage with respect to the wage income tax rate; similarly $\theta_{wL} = 2 + \varepsilon_{w,t_w} + \varepsilon_{L,t_w}$ and $\theta_L = 1 + \varepsilon_{L,t_w}$. When governments fail to account for the leakage, horizontal externalities, and second moment effects they underestimate the added volatility associated with a revenue increase and expose their revenues to unnecessary levels of risk.

$$t_w^* = \frac{\bar{w}\sigma_c^2 - \bar{y}\sigma_{w,c}}{L[\bar{c}^2\sigma_w^2 - 2\bar{c}\bar{w}\sigma_{c,w} + \sigma_c^2\bar{w}^2]} \bar{R} \quad t_w^{**} = \frac{\theta_{wL}\bar{w}\sigma_c^2 - \theta_w\bar{c}\sigma_{w,c}}{L[\theta_{wL}\bar{c}^2\sigma_w^2 - (\theta_{wL} + \theta_L)\bar{c}\bar{w}\sigma_{c,w} + \theta_{wL}\sigma_c^2\bar{w}^2]} \bar{R}$$

$$t_c^* = \frac{\bar{c}\sigma_w^2 - \bar{w}\sigma_{c,w}}{\beta[\bar{c}^2\sigma_w^2 - 2\bar{c}\bar{w}\sigma_{c,w} + \sigma_c^2\bar{w}^2]} \bar{R} \quad t_c^{**} = \frac{\theta_L\bar{c}\sigma_w^2 - \theta_{wL}\bar{w}\sigma_{w,c}}{\beta[\theta_{wL}\bar{c}^2\sigma_w^2 - (\theta_{wL} + \theta_L)\bar{c}\bar{w}\sigma_{c,w} + \theta_{wL}\sigma_c^2\bar{w}^2]} \bar{R}$$

In this special case the tax rate pairs on the minimum-variance frontier are characterized by the condition in equation (MVF.1). This condition states that the ratio of the elasticity of the variance and the elasticity of the mean of tax revenue with respect to the tax rate are equal

across tax rates. In finance, the minimum-variance frontier is often called the efficient frontier, but through this series of special cases the analysis demonstrates for the government's optimal portfolio problem the minimum-variance frontier may not be efficient.

$$(MVF.1) \quad \frac{\varepsilon_{\sigma_R^2, \tau_c}}{\varepsilon_{R, \tau_c}} = \frac{\varepsilon_{\sigma_R^2, t_w}}{\varepsilon_{R, t_w}}$$

C.2 Homogeneous Risk Attitudes Over Public and Private Consumption. In this case the government's objective function reduces to minimizing the sum of the variances of public and private consumption for a given expected level of public and private consumption. When combined, the first-order conditions produce the condition for the minimum-variance frontier in equation (MVF.2).

$$(8) \quad \min_{t_w, \tau_c} \sigma_R^2 + \sigma_c^2$$

Subject to: $\bar{R} = \tau_c \beta \bar{c} + t_w L \bar{w}$

In this case the minimum-variance frontier provides the minimum of the sum of the variances of public and private consumption for a given mean level of tax revenue. Compared to case one, the minimum-variance frontier in this case adds the elasticity of the variance of private consumption with respect to the tax rate.

$$(MVF.2) \quad \frac{\varepsilon_{\sigma_R^2, \tau_c} + \varepsilon_{\sigma_c^2, \tau_c}}{\varepsilon_{R, \tau_c}} = \frac{\varepsilon_{\sigma_R^2, t_w} + \varepsilon_{\sigma_c^2, t_w}}{\varepsilon_{R, t_w}}$$

C.3 Heterogeneous Risk Attitudes Over Public and Private Consumption. In this case the government's objective function generalizes case two, by allowing the risk attitude to differ between public and private consumption. The first-order conditions, when combined, produce the condition for the minimum-variance frontier in equation (MVF.3). In this case, the minimum-variance frontier provides the minimum welfare cost of public and private consumption for a given mean level of tax revenue.

$$(9) \quad \min_{t_w, \tau_c} - [M(\cdot, \sigma_R^2, \cdot, \cdot) + G(\cdot, \sigma_c^2)]$$

$$\text{Subject to: } \bar{R} = \tau_c \beta \bar{c} + t_w L \bar{w}$$

Compared to case two, the minimum-variance frontier in this case adds welfare weights to the elasticities in the numerator. The welfare weight on the elasticity of the variance of private consumption with respect to a tax rate is the negative of the derivative of utility with respect to the variance of private consumption multiplied by the ratio of the variance of private consumption and tax revenue, $\omega_M = -M_2 \sigma_c^2 / R$. Similarly, the welfare weight on the elasticity of the variance of public consumption with respect to a tax rate is $\omega_G = -G_2 \sigma_R^2 / R$.

$$(MVF.3) \quad \frac{\omega_M \varepsilon_{\sigma_c^2, \tau_c} + \omega_G \varepsilon_{\sigma_R^2, \tau_c}}{\varepsilon_{R, \tau_c}} - \frac{\omega_M \varepsilon_{\sigma_c^2, t_w} + \omega_G \varepsilon_{\sigma_R^2, t_w}}{\varepsilon_{R, t_w}} = 0$$

The minimum-variance frontier in this case is the same as the minimum-variance frontier in the full government's optimal portfolio problem. However, the efficient condition in the full government's problem, given in equation (4), differs from the minimum-variance frontier. Specifically, the right-hand side of the full government's condition in equation (4) is not zero because of the utility cost of deadweight loss. The optimal tax rates characterized by the condition in equation (4) tradeoff utility from tax revenue with the cost of deadweight loss and volatility. Hence, a government with the efficient tax rates could change its tax rates to decrease the welfare cost of volatility but would cause an increase in deadweight loss larger than the decrease in the cost of volatility. Therefore, in general, the efficient portfolio for the government will not be on the minimum-variance frontier.

4. EMPIRICAL METHOD

In finance, the efficient frontier is estimated using historical returns with the implicit assumption that the mean return and variance-covariance matrix are invariant to the portfolio that is held. While this is a reasonable assumption in finance because investors are small relative to the market, the previous section demonstrates it is not reasonable when applied to governments. Therefore, this section adapts finance theory to produce a method for estimating the efficient

frontier in a way that allows both the mean and variance-covariance matrix to depend on the portfolio that is held.

The returns of a portfolio can be written, as in equation 10, as the weighted sum of the returns of the possible assets. The objective is to find weights to minimize the variance of the portfolio for a given mean return. The minimum-variance frontier is found by calculating the efficient portfolio for different mean returns.

$$(10) \quad R = w_1 r_1 + w_2 r_2 + \dots = \mathbf{r}\mathbf{w}$$

The efficient weights are found by solving the first-order condition of the objective function with respect to the weights. The efficient weights, with some rearranging, are equal to the ordinary-least-squares coefficient from a regression of a constant \bar{R} on the returns of the possible assets through time. The efficient weights can be determined by a simple ordinary-least-squares regression without a constant and weighting the coefficients to sum to one.⁸

$$\min_w \sigma_R^2 = E[(\mathbf{r}\mathbf{w})^2] - (E[\mathbf{r}\mathbf{w}])^2 \quad \text{subject to } E[\mathbf{r}\mathbf{w}] = \bar{R}$$

$$2\mathbf{w}E[\mathbf{r}^2] = 2(E[\mathbf{r}]) \underbrace{(E[\mathbf{w}\mathbf{r}])}_{= \bar{R}} \quad \text{Use constraint}$$

$$\mathbf{w} = (E[\mathbf{r}^2])^{-1}(\bar{R}E[\mathbf{r}])$$

$$= \left(\frac{1}{T} \sum_{t=1}^T r_t r_t'\right)^{-1} \left(\frac{1}{T} \sum_{t=1}^T \bar{R} r_t\right) \quad \text{Sample analogue}$$

$$= (\mathbf{r}'\mathbf{r})^{-1}(\mathbf{r}'\bar{\mathbf{R}})$$

$$= \mathbf{w}^{OLS}$$

⁸The result that the efficient weights on a portfolio can be determined by a simple ordinary-least-squares regression was first shown by Britten-Jones (1999).

The ordinary-least-squares regression is biased without an intercept term; therefore, the mean of the portfolio is equal to the predicted average return, which may not equal the average return used in the regression $\mu = \hat{\mathbf{R}} \neq \bar{\mathbf{R}}$ and the variance of the portfolio is given by the residual's variance.⁹

To adapt this method to governments, the leakage, horizontal externalities, and second moment effects need to be accounted for to produce the actual mean-variance tradeoff. To account for leakage, each tax rate-base pair is considered its own asset. For example, the returns to a five percent income tax and a three percent income tax are considered separate assets because their mean and variance differs. Therefore, the assets from which the government chooses are expanded from three (income, sales, and corporate tax bases) to $3xN$ where N represents the number of different tax rates considered. However, the government is constrained to hold only one asset per tax base (e.g. the government cannot simultaneously hold a five percent income tax and a three percent income tax). The benefit of defining government assets in this way is that it allows the returns from a three percent income tax to differ from the returns of a five percent income tax in a less than proportional way, accounting for individuals' behavioral responses. The drawback of this approach is that the returns from all of these different assets are unobserved empirically and need to be estimated.

Horizontal externalities cause the returns of a given asset (e.g. a five percent income tax) to differ with the other assets held by the government (e.g. a three percent sales tax or a four percent sales tax). To account for these horizontal externalities the returns of a given asset are estimated conditional on the other assets held by the government. Therefore, to account for horizontal externalities the returns of an entire portfolio (e.g. a five percent income tax, a three percent sales tax, and a one percent corporate tax) are estimated. The benefit of estimating the returns at the portfolio level is that the horizontal externalities and second moment effects

⁹ $\mu = \bar{\mathbf{R}} - \bar{\mathbf{u}} = \bar{\mathbf{R}} - (\hat{\mathbf{R}} - \bar{\mathbf{R}}) = \hat{\mathbf{R}}$

are taken into account. The drawback is that the weights estimated using the procedure above define the mix between portfolios rather than the mix between assets.

5. APPLICATION: STATE-LEVEL MINIMUM-VARIANCE FRONTIERS

This section demonstrates the method and benefit of estimating minimum-variance frontiers. First, the data are described using simplices to depict how tax portfolios in the data have changed over time. Second, the portfolio returns are estimated using a weighted regression. Finally, the estimated portfolio returns are used to estimate minimum-variance frontiers which quantify the mean-variance tradeoff faced by governments. The analysis demonstrates the ability to compare the mean-variance tradeoff across governments and within a government across time.

5.1. Data and Basic Facts. Data from U.S. state governments is used to demonstrate the method of estimating minimum-variance frontiers. This data provides a balanced panel of fifty states through forty-eight years (1963-2010). Data on tax rates (top income, bottom income, sales, and corporate), tax revenues, and tax base characteristics are collected for all states from the Book of States and cross-checked with the Advisory Commission on Intergovernmental Relations biannual report “Significant Features in Fiscal Federalism,” and the Tax Foundation. State level economic conditions such as state level GDP and personal income are used as controls and are collected from the Bureau of Economic Analysis.

Graphing state government’s tax portfolios onto a 2-simplex demonstrates how much a government relies on each tax bases. A 2-simplex is a triangle drawn in two-space that represents three-space. In this case, the dimensions are tax revenues collected from income, sales, and corporate tax bases as a percent of the sum of these three tax bases.¹⁰ Figure 2 is an example of a 2-simplex that depicts the percent of tax revenue from income, sales, and corporate tax bases (three-space) in two-space. The nodes of the simplex denoted by A, B, and C represent tax

¹⁰The simplex is characterized by $\Delta^2 = \{(s_{income}, s_{sales}, s_{corporate}) \in R^3 | s_{income} + s_{sales} + s_{corporate} = 1\}$, where s_{income} is the percent of revenue collected from the income tax.

portfolios that rely on only one tax base.¹¹ Point A represents a tax portfolio that relies only on the sales tax, point B a tax portfolio that relies only on the income tax, and C only the corporate tax. Interior points represent mixtures between the three tax bases. Point D represents a tax portfolio that relies equally on all three tax bases. Point E represents a tax portfolio that relies fifty percent on corporate tax revenue and fifty percent on income tax revenue. Movements along the dashed lines xx, yy, and zz represent changes in the reliance of two of the three tax bases. For example moving along the dashed line zz shifts the reliance on sales and income taxes but keeps the reliance on the corporate tax fixed. Similarly, moving along the line yy shifts the reliance on the sales and income taxes but for a a tax portfolio that relies less on the corporate tax than portfolios along the line zz. Finally, moving along the line xx represents tax portfolios shifting between the income and corporate tax holding fixed the reliance on the sales tax.

Figure 3 plots the aggregate state and local tax portfolios for each year between 1951 and 2010. Between 1951 and 2010 the aggregate tax portfolio shifted away from the sales tax and toward the income tax (the horizontal-axis). In this same period, the aggregate tax portfolio shifted away from the corporate tax (the vertical-axis). Figure 4 plots each state's tax portfolio in 1955 and 2005 to demonstrate the disaggregated shift in tax portfolios. The disaggregated data in figure 4 demonstrate that a large number of states shifted their tax portfolios to rely more heavily on income taxes and less heavily on sales taxes. Despite this general trend, there are still seven states without an income tax in 2005.¹² In contrast, reliance on corporate tax revenue decreased significantly for a few states but the majority of states made only minor changes to their reliance on the corporate tax. The general trend between 1955 and 2005 was for states to become more similar in how heavily they rely on corporate taxes.

¹¹Each node of the triangle represents a portfolio made up entirely of one tax base with nodes at $(0, 0)$, $(1, 0)$, and $(.5, .866)$ corresponding to a portfolio entirely of sales, income, or corporate tax revenue respectively.

¹²States without income taxes: FL, NV, SD, TN, TX, WA, WY.

5.2. Estimating Minimum-Variance Frontiers. The first step in calculating minimum-variance frontiers is estimating the portfolio returns. Panel data is advantageous for calculating portfolio returns because it helps with the limitation that data exist for only one tax portfolio for a given state-year observation (i.e., the tax portfolio that actually existed for that state-year). With panel data additional portfolios can be formed by appropriately weighting observations from other states according to the inverse probability weights. Inverse probability weights calculate the probability that any state-year observation could have been observed in a given state, based on the characteristics of the state-year observation. For example, the probability that observations in Wisconsin could have been observed in Minnesota is higher than the probability that observations in Wisconsin could have been observed in California. The inverse probability weights are calculated using a probit model with a dependent variable described by an indicator function equal to one if the observation occurred in a given state, and tax rate and economic variables as the independent variables. A separate probit is run for each state calculating fifty weights, one for each state, for each state-year observation.

Tax revenue returns for different portfolios are calculated using coefficients estimated from the weighted regression in 11 where β_3 and β_4 are vectors of coefficients for the tax and economic variables respectively and $\mathbf{T}_t\beta_1$ is a vector of time trend variables.¹³ In principal, an unlimited number of portfolio returns can be calculated from this regression by substituting different sets of tax rates. In practice, the number of portfolios depend on how many values a given tax rate is allowed to take. If the number of different values is constant across the four different tax rates the number of portfolios is given by 4^n where n is the number of values each tax rate can take. For this application each tax rate is allowed to take on ten different values that are evenly spaced between zero and the maximum tax rate observed in the data, producing 1,048,576 different portfolios. The regression method calculates the optimal mix of the calculated portfolio returns

¹³Including the time trend variables is equivalent to detrending the variables with respect to time.

estimating a continuous minimum-variance frontier from discrete choices of portfolios.

$$(11) \quad \log(R_{i,t}) = \beta_0 + \mathbf{t}_t\beta_1 + \log(\boldsymbol{\tau}_{i,t})\beta_3 + \log(\mathbf{x}_{i,t})\beta_4 + \varepsilon$$

5.3. Analysis of Minimum-Variance Frontiers. Figure 5 graphs the estimated minimum-variance frontiers for Idaho (solid line) and Nevada (dashed line). Idaho's minimum-variance frontier is considerably steeper than Nevada's implying Idaho's tax base is less volatile than Nevada's. Estimating minimum-variance frontiers for state governments is useful for understanding an individual state's mean-variance tradeoff and for comparing mean-variance tradeoffs across states. This comparison is particularly useful in considering the costs to different states due to changes in federal policies. For example, consider the costs to state governments from a decrease in intergovernmental transfers from the federal government. Even if this decrease was proportional across states the costs may not be because the increase in volatility caused by states responses differs across states. In this example, Nevada faces larger costs in terms of volatility than Idaho does due to the increase in tax revenue collections.

Figure 6 graphs California's estimated minimum-variance frontier and the actual portfolios (open circles) in the past 48 years. The horizontal distance between a portfolio and the efficient frontier is the additional variance the portfolio has relative to a portfolio with the same mean on the minimum-variance frontier. Similarly, the vertical distance between a portfolio and the minimum-variance frontier is how much less mean revenue the portfolio has relative to a portfolio with the same variance on the minimum-variance frontier.

Through time California has increased both the mean and variance of the tax revenues it collects. The arch formed by the actual portfolios held by California in the past 48 years is flatter than the estimated minimum-variance frontier. Hence, California has exposed its revenues to inefficient levels of risk, quantified by the distance between the actual portfolio and the minimum-variance frontier. Although other considerations of optimal taxation may cause state governments

to choose portfolios off of the minimum-variance frontier, comparing these portfolios with the estimated minimum-variance frontier quantifies the cost in terms of additional volatility (or lower mean) from choosing such a portfolio.

This paper focuses on the tradeoff between volatility and deadweight which causes governments to efficiently choose portfolios that are not on their minimum-variance frontier. In general, tradeoffs involving redistribution may also cause governments to be off of their minimum-variance frontier, which may help explain the difference between California's minimum-variance frontier and actual portfolios. However, volatility is an additional cost state governments need to consider when making their decisions on redistribution.

FIGURE 2. Simplex Example.

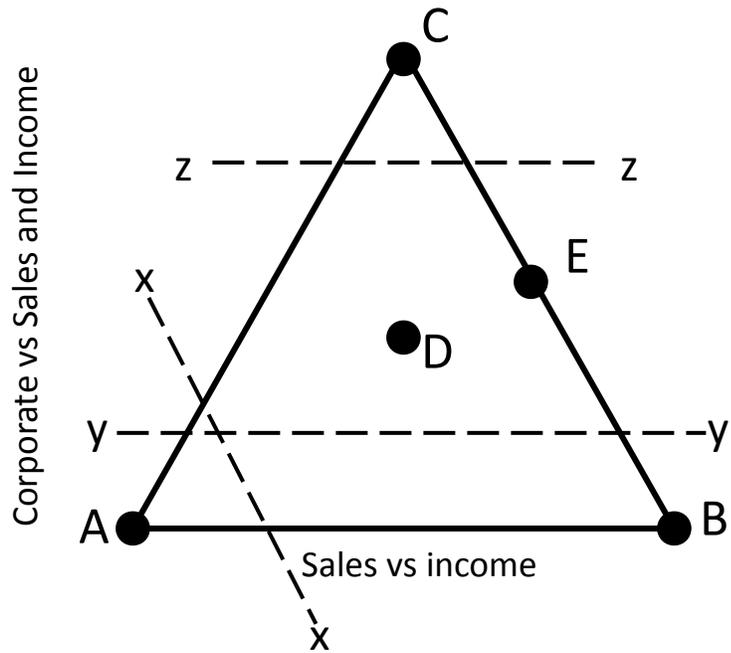


FIGURE 3. Aggregate State Tax Portfolios Over Time.

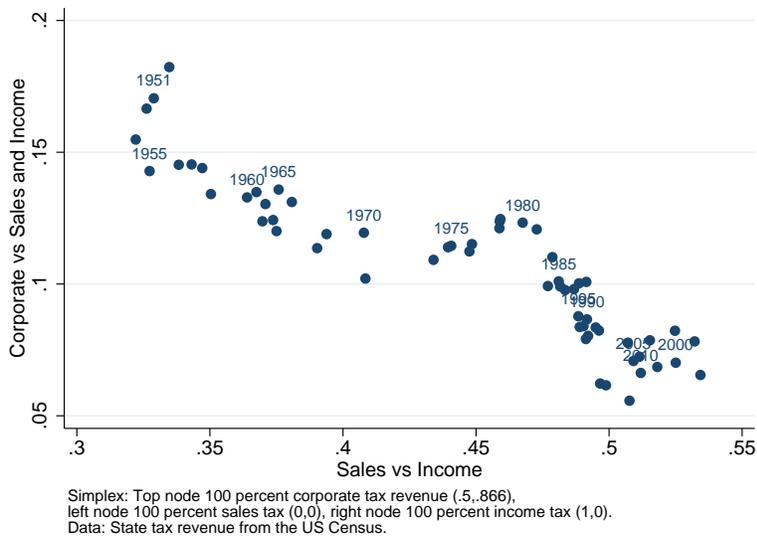
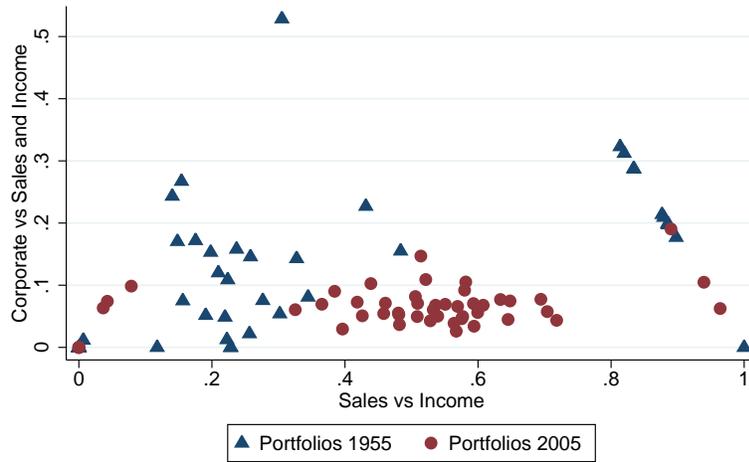


FIGURE 4. State Tax Portfolios 1955 and 2005.



Simplex: Top node 100 percent corporate tax revenue (.5,.866),
 left node 100 percent sales tax (0,0), right node 100 percent income tax (1,0).
 Alaska and New Hampshire not shown.
 Data: State tax revenue from the US Census.

FIGURE 5. Idaho and Nevada Estimated Efficient Frontiers.

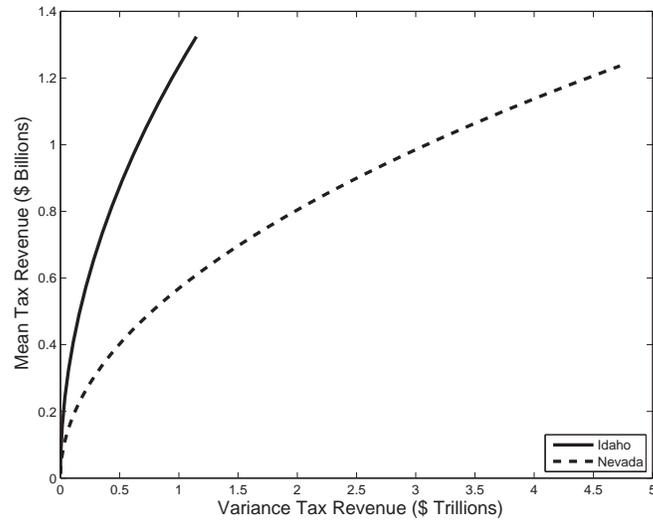
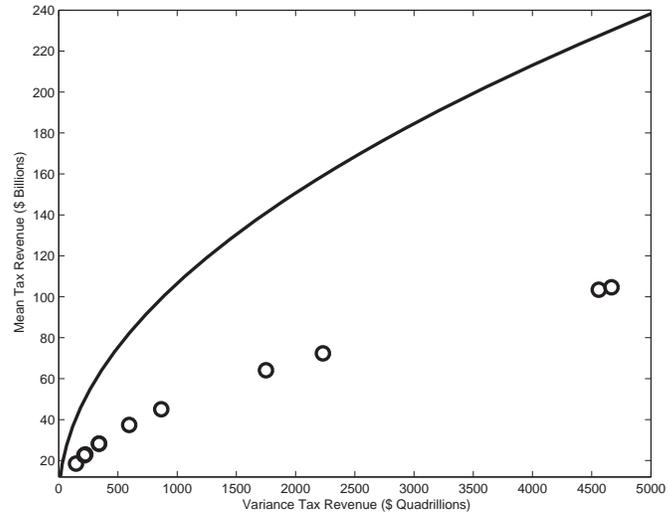


FIGURE 6. California Estimated Efficient Frontier and Actual Portfolios.



6. CONCLUSION

The economics literature has long understood that states have “few if any of the unique powers that make fluctuating tax yields a matter of minor concern to the federal government” (Groves and Kahn 1952). This paper considers the desire of stable tax revenues for state and local governments through an optimal portfolio analysis within a utility framework. The optimal portfolio analysis demonstrates the ability of governments to hedge idiosyncratic risk involved with a given tax base. Traditional portfolio analysis is adapted to account for the unique position of a government as a large player. The utility framework demonstrates the tradeoffs faced by governments between volatility and deadweight loss and between public and private consumption volatility. Therefore, in general the government’s objective is to produce the optimal level of tax revenue volatility, but not necessarily to minimize tax revenue volatility.

This paper develops a method for estimating the minimum-variance frontier for governments, which have assets that respond to the weight placed on them. Data from U.S. state governments are used to demonstrate the method, providing a comparison of state tax portfolios across states and historically within states. This analysis quantifies the cost in terms of forgone mean levels of revenue (or additional variance) of the portfolios state governments have held historically. Groves and Kahn in their 1952 paper discuss the tradeoff between stability in tax revenue and the desire for redistribution of wealth. With the method from this paper of estimating government minimum-variance frontiers governments can quantify the cost of additional volatility caused by their redistribution policies which is necessary for them to make an informed decision on what level of redistribution to implement.

REFERENCES

- [1] J. Albrecht, A. Bjorklund, and S. Vroman. Is there a glass ceiling in sweden? *Journal of Labor Economics*, 21(1), 2003.
- [2] M.G. Allingham and A. Sandmo. Income tax evasion: A theoretical analysis. *Journal of public economics*, 1(3-4):323–338, 1972.
- [3] A.S. Blinder. Wage discrimination: reduced form and structural estimates. *Journal of Human resources*, pages 436–455, 1973.
- [4] N. Bloom. The impact of uncertainty shocks. *Econometrica*, 77(3):623–685, 2009.
- [5] N. Bloom, M. Floetotto, and N. Jaimovich. Really uncertain business cycles. 2009.
- [6] M. Britten-Jones. The Sampling Error in Estimates of Mean-Variance Efficient Portfolio Weights. *The Journal of Finance*, 54(2):655–671, 1999.
- [7] D. Bruce and W.F. Fox. E-commerce in the context of declining state sales tax bases. *National Tax Journal*, 53(4; PART 3):1373–1388, 2000.
- [8] D. Bruce, W.F. Fox, and M.H. Tuttle. Tax Base Elasticities: A Multi-State Analysis of Long-Run and Short-Run Dynamics. *Southern Economic Journal*, 73(2):315–341, 2006.
- [9] M.A. Desai. The corporate profit base, tax sheltering activity, and the changing nature of employee compensation, 2002.
- [10] D.A. Dickey and W.A. Fuller. Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica: Journal of the Econometric Society*, pages 1057–1072, 1981.
- [11] J. DiNardo, N.M. Fortin, and T. Lemieux. Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach. *Econometrica*, 64(5):1001–1044, 1996.
- [12] R.F. Dye and T.J. McGuire. Growth and variability of state individual income and general sales taxes. *National Tax Journal*, 44(1):55–66, 1991.
- [13] W.F. Fox and C. Campbell. Stability of the state sales tax income elasticity. *National Tax Journal*, 37(2):201–212, 1984.
- [14] S. Gilchrist, J.W. Sim, and E. Zakrajšek. Uncertainty, financial frictions, and investment dynamics. Technical report, 2010.
- [15] R.H. Gordon and J.B. Slemrod. Are "real" responses to taxes simply income shifting between corporate and personal tax bases?, 1998.
- [16] P. Gottschalk and R. Moffitt. The rising instability of US earnings. *The Journal of Economic Perspectives*, 23(4):3–24, 2009.

- [17] C.W.J. Granger and P. Newbold. Spurious regressions in econometrics. *Journal of Econometrics*, 2(2):111–120, 1974.
- [18] J.R. Green and E. Sheshinski. Approximating the efficiency gain of tax reforms. *Journal of Public Economics*, 11(2):179–195, 1979.
- [19] H.M. Groves and C.H. Kahn. The stability of state and local tax yields. *The American Economic Review*, 42(1):87–102, 1952.
- [20] A.C. Harberger. The measurement of waste. *The American Economic Review*, 54(3):58–76, 1964.
- [21] J.R. Hines. State Fiscal Policies and Transitory Income Fluctuations. *Brookings Papers on Economic Activity*, 2010(2):313–337, 2011.
- [22] J.R. Hines Jr. Michigan flirtation with the Single Business Tax. *Ann Arbor*, 1001:48109–1234, 2002.
- [23] K. Hirano, G.W. Imbens, and G. Ridder. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189, 2003.
- [24] P.W. Holland. Statistics and causal inference. *Journal of the American Statistical Association*, pages 945–960, 1986.
- [25] P. Kline. Blinder-oaxaca as a reweighting estimator, 2010.
- [26] S. Ludvigson and C. Steindel. How important is the stock market effect on consumption? *Economic Policy Review*, (Jul):29–51, 1999.
- [27] J.A.F. Machado and J. Mata. Counterfactual decomposition of changes in wage distributions using quantile regression. *Journal of applied Econometrics*, 20(4):445–465, 2005.
- [28] B. Melly. Decomposition of differences in distribution using quantile regression. *Labour Economics*, 12(4):577–590, 2005.
- [29] R.C. Merton. An analytic derivation of the efficient portfolio frontier. *Journal of financial and quantitative analysis*, 7(04):1851–1872, 1972.
- [30] J.A. Mirrlees. An exploration in the theory of optimum income taxation. *The Review of Economic Studies*, 38(2):175–208, 1971.
- [31] J. Mossin. Risk and uncertainty. *New York*, 1969.
- [32] R. Oaxaca. Male-female wage differentials in urban labor markets. *International economic review*, 14(3):693–709, 1973.
- [33] Y. Otsuka and B.M. Braun. The random coefficient approach for estimating tax revenue stability and growth. *Public Finance Review*, 27(6):665, 1999.

- [34] M.O. Ravn and H. Uhlig. On adjusting the hodrick-prescott filter for the frequency of observations. *Review of Economics and Statistics*, 84(2):371–376, 2002.
- [35] R.S. Sobel and R.G. Holcombe. Measuring the growth and variability of tax bases over the business cycle. *National Tax Journal*, 49:535–552, 1996.
- [36] N. Stern. From the static to the dynamic: some problems in the theory of taxation. *Journal of public economics*, 47(2):273–297, 1992.
- [37] J.E. Stiglitz. A re-examination of the modigliani-miller theorem. *The American Economic Review*, 59(5):784–793, 1969.
- [38] J.E. Stiglitz and P. Dasgupta. Differential taxation, public goods, and economic efficiency. *The Review of Economic Studies*, 38(2):151–174, 1971.
- [39] J.H. Stock and M.W. Watson. A simple estimator of cointegrating vectors in higher order integrated systems. *Econometrica*, 61(4):783–820, 1993.
- [40] H.R. Varian. Redistributive taxation as social insurance. *Journal of Public Economics*, 14(1):49–68, 1980.
- [41] J. Vavra. Inflation dynamics and time-varying uncertainty: New evidence and an ss interpretation.

7. APPENDIX: HIGHER MOMENT DECOMPOSITION

In the text the expected utility is assumed to be fully characterized by the first two moments of the public and private good, which is sufficient when the goods are jointly normally distributed or when the utility function is quadratic. In general, the expected utility can be written as a function of all higher moments Ω . If the utility function is additively separable, such that $U_{1,4} = 0$, then the level of social welfare can be written as the second line in the equation below.

$$\int U(c, \beta, L, G) f(\bar{c}, \bar{R}, \Omega) \equiv M((\bar{c}, \beta, L, \bar{R}, \Omega)$$

$$= M(\bar{c}, \Omega_c, \beta, L) + G(\bar{R}, \Omega_R) \quad \text{When } U_{1,4} = 0$$

To demonstrate this transformation consider a Cobb Douglas utility where total consumption is assumed to be distributed uniformly with mean μ and standard deviation σ . Writing the utility function in terms of total consumption c and β gives the following form where the density function is $\frac{1}{2\sigma\sqrt{3}}$ for $c \in [-\sqrt{3}\sigma, \sqrt{3}\sigma]$ and zero everywhere else.

$$\begin{aligned}
E[U(c, \beta)] &= E[\log c + \alpha \log \beta + (1 - \alpha) \log(1 - \beta)] \\
&= E[\log[c]] + \alpha \log \beta + (1 - \alpha) \log(1 - \beta) \\
&= \int_{-\sqrt{3}\sigma}^{\sqrt{3}\sigma} \log c \frac{1}{2\sigma\sqrt{3}} dc + \alpha \log \beta + (1 - \alpha) \log(1 - \beta) \\
M(\mu, \sigma^2, \beta) &= \frac{(\sigma\sqrt{3} + \mu)(\log(\mu + \sigma\sqrt{3}) - 1) + (\sigma\sqrt{3} - \mu)(\log(\mu - \sigma\sqrt{3}) - 1)}{2\sigma\sqrt{3}} \\
&\quad + \alpha \log \beta + (1 - \alpha) \log(1 - \beta)
\end{aligned}$$

The preceding line is a function of the mean, standard deviation, and β alone.

8. APPENDIX: CONSUMPTION BASE DECOMPOSITION

In the text private consumption of the representative agent is decomposed into consumption that is taxed and consumption that is untaxed such that the fraction β of total consumption is taxed and $(1 - \beta)$ is untaxed. This decomposition changes the variables from two consumption goods into total consumption and the fraction spent on taxable items. This section demonstrates the change of variables and its benefits.

First, start with two goods B, N such that the consumption of B is taxed and the consumption of N is not taxed and the representative agent has utility $\mathbb{U}(B, N)$ over the two goods. By definition $B = \beta c$ and $N = (1 - \beta)c$. The utility function can be written as a function of β and c by substituting these equations in for B and N .¹⁴ The budget constraint is given below written both as a function of B and N and β and c .

$$\begin{aligned} W &= (1 + \tau_c)B + N \\ &= (1 + \tau_c)\beta c + (1 - \beta)c \\ &= c(1 + \beta\tau_c) \end{aligned}$$

Now we want to know the welfare impact of a tax change. We can separate the impact into the income effect and the substitution effect where the substitution effect is the deadweight loss from the behavioral responses.

$$\begin{aligned} \frac{\partial \mathbb{U}(B, N)}{\partial \tau_c} &= \mathbb{U}_1 \frac{\partial B}{\partial \tau_c} + \mathbb{U}_2 \frac{\partial N}{\partial \tau_c} \\ &= \mathbb{U}_1 \left(S_{B, \tau_c} - \frac{\partial B}{\partial W} B \right) + \mathbb{U}_2 \left(S_{N, \tau_c} - \frac{\partial N}{\partial W} B \right) \quad \text{Slutsky Decomposition} \\ &= \underbrace{\mathbb{U}_1 S_{B, \tau_c} + \mathbb{U}_2 S_{N, \tau_c}}_{\text{Substitution Effect}} - \underbrace{\left(\mathbb{U}_1 \frac{\partial B}{\partial W} B + \mathbb{U}_2 \frac{\partial N}{\partial W} B \right)}_{\text{Income Effect}} \end{aligned}$$

¹⁴If the utility function is homothetic then the utility function can be written as $\mathbb{U}(B, N) = v(\beta)U(c)$ otherwise $\mathbb{U}(B, N) = U(c, \beta)$.

The benefit of writing the utility in terms of β and c is that $U_1 \frac{\partial c}{\partial \tau_c}$ captures the income effect and $U_2 \frac{\partial \beta}{\partial \tau_c}$ captures the behavioral response and deadweight loss.

$$\begin{aligned}
 - \underbrace{\left(\mathbb{U}_1 \frac{\partial B}{\partial W} B + \mathbb{U}_2 \frac{\partial N}{\partial W} B \right)}_{\text{Income Effect}} &= \mathbb{U}_1 \frac{B}{c} \frac{\partial c}{\partial \tau_c} + \mathbb{U}_2 \frac{B(1-\beta)}{c\beta} \frac{\partial c}{\partial \tau_c} \\
 &= \mathbb{U}_1 \beta \frac{\partial c}{\partial \tau_c} + \mathbb{U}_2 (1-\beta) \frac{\partial c}{\partial \tau_c} \\
 &= U_1 \frac{c}{\tau_c}
 \end{aligned}$$

The first equality holds because of the following.

$$\begin{aligned}
 \frac{\partial B}{\partial W} &= \frac{\partial \beta c}{\partial W} \\
 &= \beta \frac{\partial c}{\partial W} \\
 &= \frac{\beta}{1 + \tau_c \beta} && \text{where } c = \frac{W}{1 + \tau_c \beta} \\
 &= -\frac{\partial c}{\partial \tau_c} \frac{1}{c} && \text{where } \frac{\partial c}{\partial \tau_c} = -\frac{W\beta}{(1 + \tau_c \beta)^2} = -\frac{c\beta}{(1 + \tau_c \beta)}
 \end{aligned}$$

$$\frac{\partial N}{\partial W} = \frac{\partial(1 - \beta)c}{\partial W}$$

$$= (1 - \beta) \frac{\partial c}{\partial W}$$

$$= \frac{(1 - \beta)}{1 + \tau_c \beta}$$

$$\text{where } c = \frac{W}{1 + \tau_c \beta}$$

$$= -\frac{\partial c}{\partial \tau_c} \frac{(1 - \beta)}{\beta c}$$

$$\text{where } \frac{\partial c}{\partial \tau_c} = -\frac{W\beta}{(1 + \tau_c \beta)^2} = -\frac{c\beta}{(1 + \tau_c \beta)}$$

The last equality holds because of the following.

$$U_1 = \mathbb{U}_1 \frac{\partial B}{\partial c} + \mathbb{U}_2 \frac{\partial N}{\partial c} = \mathbb{U}_1 \beta + \mathbb{U}_2 (1 - \beta)$$

Now show the deadweight loss calculation.

$$\underbrace{\mathbb{U}_1 S_{B,\tau_c} + \mathbb{U}_2 S_{N,\tau_c}}_{\text{Substitution Effect}} = 0$$

$$\begin{aligned} U_2 \frac{\partial \beta}{\partial \tau_c} &= \frac{\partial \beta}{\partial \tau_c} \left(\mathbb{U}_1 \frac{\partial B}{\partial \beta} + \mathbb{U}_2 \frac{\partial N}{\partial \beta} \right) \\ &= \frac{\partial \beta}{\partial \tau_c} \left(\mathbb{U}_1 \frac{c}{1 + \beta \tau_c} - \mathbb{U}_2 \frac{c(1 + \tau_c)}{1 + \beta \tau_c} \right) \\ &= \frac{\partial \beta}{\partial \tau_c} \left(\mathbb{U}_1 \left[\frac{c}{1 + \beta \tau_c} - \frac{\mathbb{U}_2 c(1 + \tau_c)}{\mathbb{U}_1 (1 + \beta \tau_c)} \right] \right) \\ &= \frac{\partial \beta}{\partial \tau_c} \left(\mathbb{U}_1 \left[\frac{c}{1 + \beta \tau_c} - \frac{c}{1 + \beta \tau_c} \right] \right) \\ &= 0 \end{aligned}$$

where from totally differentiating the budget constraint $\frac{\partial B}{\partial \beta} = \frac{c}{1 + \beta \tau_c}$, $\frac{\partial N}{\partial \beta} = -\frac{c(1 + \tau_c)}{1 + \beta \tau_c}$, and from the individual's optimization $\frac{\mathbb{U}_2}{\mathbb{U}_1} = \frac{1}{1 + \tau_c}$.

$$\begin{aligned} \frac{\partial B}{\partial \beta} &= c + \beta \frac{\partial c}{\partial \beta} & \text{Total Differentiate B.C. } 0 &= (1 + \tau_c \beta)dc + \tau_c c d\beta \\ &= c - c \frac{\beta \tau_c}{1 + \beta \tau_c} & dc/d\beta &= -\tau_c c / (1 + \tau_c \beta) \\ &= \frac{c}{1 + \beta \tau_c} \end{aligned}$$

$$\begin{aligned}\frac{\partial N}{\partial \beta} &= -c + (1 - \beta) \frac{c}{\beta} \\ &= -c - c \frac{(1 - \beta)\tau_c}{1 + \beta\tau_c} \\ &= -\frac{c(1 + \tau_c)}{1 + \beta\tau_c}\end{aligned}$$

E-mail address: seegert@umich.edu