State Tax Risk

Sebastien Bradley* Da Huang † Nathan Seegert ‡

February 27, 2025

Abstract

State tax policy distorts individuals' decisions on entrepreneurship, investment, and effort. We show that different tax policies have different levels of risk based on the timing of when taxes are due. We use a novel measure of *tax risk* based on an Arrow-Debreu framework, which captures how burdensome a tax series is by adjusting up (down) the tax liabilities in a poor (rich) state of the world through the pricing kernels. Therefore, *tax risk* is based on the covariance (between taxes and states of the world) rather than the variance, which is typically studied. Our approach encompasses popular considerations in tax design, including surprise, asymmetry, and state-dependent characteristics. The *tax risk* varies greatly across states which can potentially affect real economic outcomes on the state level.

Keywords: State taxation, Exposure to risk
JEL Classification: H2, H7, R51, E2

*School of Economics, LeBow College of Business, Drexel University. 1023 Gerri C. LeBow Hall, 3220 Market Street, Philadelphia, PA 19104. Office phone: (215) 571-4193. E-mail: sbradley@drexel.edu

†Department of Finance, D'Amore-McKim School of Business, Northeastern University, 360 Huntington Avenue, Boston, MA 02115, US. E-mail: da.huang@northeastern.edu

‡*Corresponding author:* Department of Finance, University of Utah, 1655 East Campus Center Drive, Salt Lake City, UT 84112, US. Office phone: (801) 585-7131. E-mail: nathan.seegert@eccles.utah.edu

1 Introduction

Individuals routinely balance risk and reward when making investment and effort decisions. The system of federal and state taxes distorts this tradeoff in two important ways. On the one hand, taxes reduce the net returns to investment and effort, dampening both. Estimates suggest that for every $100 raised, the economy bears approximately $30 in lost productive activity. On the other hand, the tax system alters the risk profile of these decisions. As early as Domar and Musgrave (1944), the literature has noted that taxes can function as a form of insurance (Blundell, Pistaferri, and Preston, 2008; Splinter, 2012; Dauchy, Navarro-Sanchez, and Seegert, 2021). Taxation effectively makes the government a stakeholder in individual entrepreneurial ventures.

The tax code distorts risk differentially by state. One distortion comes from states taxing income and sales differently. Another distortion comes from the federal income tax, which, while administered uniformly, interacts with states differentially based on differences in state-level income distributions. The tax code, therefore, affects the balance of risk and reward across states differentially.

We develop a novel method for quantifying *tax risk* to quantify the differences in risk across states. The macro-finance literature documents that income taxes can mitigate the transmission of macroeconomic shocks, but a consensus measure of the insurance and risk value provided by the tax remains elusive (Kaplan and Violante, 2010; Larrimore, Mortenson, and Splinter, 2015; Armour, Burkhauser, and Larrimore, 2016; Heathcote, Perri, Violante, and Zhang, 2023). We fill this gap with a measure that encapsulates key elements of tax uncertainty while remaining straightforward to implement. The core premise of our approach is to view tax payments as negative returns on assets. In doing so, we price streams of tax payments using a stochastic discount factor framework. Specifically, we calculate *tax risk* as the weighted average of tax payments, with weights determined by a pricing kernel. Although our method allows for general forms of the pricing kernel, the Arrow-Debreu model offers a natural framework in which taxpayers' intertemporal marginal rate of substitution or risk

preferences weigh tax payments. Importantly, we net out the average level of tax payments so that the asset has a zero expected return. This adjustment isolates the risk effect—capturing the covariance between tax payments and the state of the world—without conflating it with average tax levels.

Our measure possesses several attractive features. First, it carries an inherent economic interpretation: it represents the additional amount an individual would be willing to pay over the certainty equivalent to avoid a stream of time-varying tax obligations. Second, it captures potential asymmetries in risk stemming from the diminishing marginal utility of income across different economic states. Third, the measure is easy to implement and flexible enough to be applied to a broad range of contexts involving streams of obligations or receipts. By contrast, traditional statistical measures of tax risk based on intertemporal variances or standard deviations—although simple to compute—fail to distinguish between volatility and growth components clearly and do not provide direct insight into behavioral ramifications. While methods such as detrending (e.g., via Hodrick-Prescott filtering) have been proposed to separate these components, they require additional assumptions and often function as "black box" techniques without addressing the asymmetry inherent in risk.

To illustrate our approach, consider an example with two states of the world—good times and bad times—each occurring with a 50% probability. Suppose that in good times, the marginal utility of consumption is 0.9, and in bad times, it is 1.1, so that the pricing kernel assigns weights of 0.45 and 0.55, ($0.5 * 0.9 = 0.45$ in good times and $0.5 \times 1.1 = 0.55$ in bad times). In jurisdiction A, tax payments are \$90 in good times and \$110 in bad times, whereas in jurisdiction B, they are \$120 and \$80, respectively. Although both jurisdictions have an average tax payment of \$100, the standard deviation and the correlation with the state of the world differ. The standard deviation of tax payments is greater in jurisdiction B (\$20) than in A (\$10). Tax payments are perfectly negatively correlated in jurisdiction A and perfectly positively correlated in jurisdiction B. The *tax risk price* is the difference between the total price and certainty equivalent price. The total price is the weighted

average of the tax payments; $-\$101 = -\$90 \times 0.45 + -\$110 \times 0.55$ in jurisdiction A and $-\$98 = -\$120 \times 0.45 + -\$80 \times 0.55$ in jurisdiction B. The certainty equivalent price is $-\$100 = -\$100 \times 0.45 + -\$100 \times 0.55$. Therefore, the *tax risk price* is -$1 = -$101 - (-$100) in jurisdiction A and $2= -$98 - (-$100) in jurisdiction B. This means that people have a negative value of these payments relative to the certainty equivalent (meaning these are risky payments) in jurisdiction A and they have a positive value of these payments relative to the certainty equivalent (meaning these payments provide insurance) in jurisdiction B. Despite having a larger standard deviation of tax payments in jurisdiction B, the tax payments in jurisdiction A are riskier than those in B precisely because of how tax payments are correlated with the different states of the world.

The conceptual framework underlying our method generalizes easily. The Arrow-Debreu pricing model aggregates any series of payouts into a single measure that captures risk. Its implementation requires only a column of data representing tax payments, a corresponding measure of the state of the world (such as consumption, housing prices, or government revenues), and two calibrated parameters—a discount rate (typically around 0.98) and a risk preference parameter (often set between 1 and 10, Chetty (2006)). Moreover, our method can be seamlessly integrated with simulation models like NBER's TAXSIM or property tax simulators (Bradley, Huang, and Seegert, 2023) to isolate key drivers of tax risk.

We apply our methodology in two main settings that underscore the tax risk (or insurance) across states. First, using panel data from the IRS's Statistics of Income for 1979–1990, we calculate Tax insurance for the federal income tax. Our findings suggest that the average household in the 1980s would have been willing to pay approximately $738 annually for the insurance provided by the federal income tax. Importantly, we also document significant heterogeneity across states. For instance, states such as Wyoming, Montana, and Colorado exhibit insurance values exceeding $1,400 per year, while states like Vermont, Delaware, and New Mexico offer less than $360. These differences are not solely due to variations in average tax amounts but are also driven by how income shocks are distributed across states.

Second, we extend our analysis to state income and sales taxes using data from the Panel Study of Income Dynamics, the Consumer Expenditure Survey (1967–2010), and NBER's TAXSIM program. This application reveals considerable heterogeneity in the insurance value provided by state-level tax systems. As expected, states with markedly different tax regimes—for example, states with no income tax such as Alaska and Florida versus states with high top rates like California—exhibit divergent levels of Tax insurance. Our measure quantifies this heterogeneity in economic terms, showing that the insurance value ranges from over $100 in states like Massachusetts and Connecticut to nearly zero in states such as Arizona and Montana. Moreover, our results indicate that neither income taxes nor sales taxes are universally riskier; instead, the relative riskiness depends critically on the specifics of the tax base, including the progressivity of the income tax and the treatment of certain goods and services in the sales tax base.

The implications of our findings extend well beyond the immediate welfare consequences of *tax risk*. The degree of risk and insurance provided by both federal and state tax systems will likely affect broader economic outcomes, including investment, entrepreneurial effort, and CEO compensation (Hassett and Metcalf, 1999; Bloom, Bond, and Van Reenen, 2007; Graham, Hanlon, Shevlin, and Shroff, 2014; Jacob, Wentland, and Wentland, 2022). Building on previous studies that link tax risk to mergers and acquisitions, CEO pay structures, and precautionary cash holdings, our work offers a quantitative framework to assess the welfare value of tax systems. Furthermore, differential tax insurance across jurisdictions may influence residents' risk-taking behavior, potentially contributing to disparities in income, wealth, and overall economic opportunity.

In sum, this paper introduces an innovative measure of Tax insurance that not only captures the level and timing of tax liabilities but also offers a direct, welfare-relevant interpretation of tax risk. By bridging gaps in the existing literature and highlighting critical considerations for state and local taxation policies, our approach provides new insights into the trade-offs inherent in tax policy design and its far-reaching implications for economic

behavior.

The remainder of the paper proceeds as follows. Section 2 provides background about the federal income tax related to risk and insurance. Section 3 defines our measure of *Tax insurance* and discusses its properties. Section 4 describes our data for the application and section 5 provides our estimates. Section 7 concludes.

2 Background

A core question in household finance revolves around the degree to which income shocks are transmitted to consumption (Blundell et al., 2008; Kaplan and Violante, 2010, 2014). The literature has focused on the extent to which income shocks are transitory or permanent (MaCurdy, 1982; Abowd and Card, 1989; Autor, Katz, and Kearney, 2008; Meghir and Pistaferri, 2004) and the relative ability of households to insure against permanent and transitory income shocks (Cutler and Katz, 1992; Altonji and Siow, 1987; Attanasio and Weber, 1995; Aguiar and Bils, 2015; Attanasio and Pavoni, 2011; Attanasio, Hurst, and Pistaferri, 2014). For example, there is a large literature considering the permanent income hypothesis (PIH), buffer-stock savings, impatient consumers, and models with liquidity constraints (Zeldes, 1989; Carroll and Kimball, 2001).

The federal individual income tax system plays a key role in enabling households to smooth income shocks, with the result that after-tax income is at least 15% smoother than before-tax income (Heathcote, Perri, and Violante, 2010; Debacker, Heim, Panousi, Ramnath, and Vidangos, 2013; Dauchy et al., 2021). This form of insurance stems from the fact that households' tax liability under the federal income tax system is pro-cyclical and tax rates are progressive. Over the last 30 years, the amount of insurance that the federal income tax provides against income shocks has decreased (Auerbach and Feenberg, 2000; Pechman, 1973; Bargain, Dolls, Immervoll, Neumann, Peichl, Pestel, and Siegloch, 2015; Dauchy et al., 2021).

We build on this literature to consider the difference in insurance provided by the federal income tax across states. At first, it might seem that the federal income tax would provide the same amount of insurance across states since the tax schedule does not differ by state. However, this neglects to consider key differences across states that interact with the federal income tax that can cause insurance amounts to differ. Dauchy et al. (2021) highlight the fact that the federal income tax provides different amounts of insurance at different income levels. Therefore, states with different income distributions will provide different insurance levels. We build on this first insight by drawing on the literature about permanent and transitory income shocks. To the extent that states experience temporary and permanent income shocks of differing magnitudes and that the federal income tax provides different levels of insurance in response to these shocks, then the level of insurance will differ across states.

In this section, we quantify households' *Tax insurance* from federal income taxation in order to demonstrate the applicability of our concept of *Tax insurance* to topics at the intersection of macro-labor and public finance. Differences across states in *Tax insurance* resulting from households' federal income tax payments constitute insurance rather than risk.

3 Framework

The key insight to tax insurance is that individuals care about the timing of income tax in addition to its level or variability. Tax schedules that generate the same expected tax liability impose different burdens on households with different timing of those payments. In particular, the schedule requiring relatively higher tax payments during bad economic times is more burdensome. In contrast, progressivity in the federal income tax provides insurance. Previous studies focus on how after-tax income has a smaller variance (or less volatility) than pre-tax income. We add to this result and argue that risk and insurance also depend on when

the payments occur relative to the state of the world. The value of $100 can be very different in a recession versus a boom. It can also be very different for an 18-year-old starting out their career or a 50-year-old who is well established in theirs. These two examples highlight that the federal income tax can provide both economic-cycle insurance in the short run and life-cycle insurance over decades. The key insight is straightforward – the same tax liability is more burdensome when the individual has less means to pay.

We develop a *Tax Insurance* measure to capture this insight. We first calculate a weighted sum of tax payments, up-weighted when the taxpayer is poor and down-weighted when the taxpayer is rich. We then calculate a certainty equivalent tax payments. *Tax Insurance* is the difference between the two. The key is the weights used in calculating the weighted sum. A constant weight, suggesting that the taxpayer is equally rich at all times, results in a weighted sum of tax payments equal to the certainty equivalent. In this case, *Tax Insurance* equals zero because taxes do not provide insurance value as the taxpayer does not care at all when to pay taxes. If the weight is low (i.e., the taxpayer is rich) when the tax payment is high, *Tax Insurance* is positive, and vice versa.

In the remainder of this section, we formalize the framework for pricing *Tax Insurance*. We then discuss our measure's connection to, and advantage over, other commonly used risk and insurance measures.

3.1 Defining Tax Insurance

We first evaluate each tax payment $x_{i,t}$ for taxpayer i at time t. The tax burden in terms of the utility of this payment depends on the taxpayer's past and current consumption, $c_{i,t-1}$ and $c_{i,t}$, and their utility function $u(\cdot)$. Specifically, the marginal rate of substitution in the Euler equation, $u^{'}(c_{i,t})/u^{'}(c_{i,t-1})$ provides a weight to the tax payment $x_{i,t}$ to quantify the burden of this payment. This weight is exactly the weight from an Arrow-Debreu security.

We then can define the price of this payment, or the *tax burden*, as

$$p_{i,t} = \frac{u^{'}(c_{i,t})}{u^{'}(c_{i,t-1})} x_{i,t}. \tag{1}$$

For any risk-averse taxpayer with twice-differentiable utility ($u'(c) > 0$ and $u''(c) < 0$), their tax burden, $p_{i,t}$, will exceed their tax payment, $x_{i,t}$, when the state of the world is bad, i.e., $c_{i,t} < c_{i,t-1}$, and vice versa when it is good. More generally, taxpayer i uses $u^{'}(c_{i,t})/u^{'}(c_{i,t-1})$ as their pricing kernel, denoted as $k_{i,t}$. Her total tax burden is

$$\mathbb{P}_i = \frac{1}{T} \sum_{t=1}^{T} k_{i,t} x_{i,t}. \tag{2}$$

We decompose this tax burden into the certainty equivalent and *Tax Insurance*. The certainty equivalent is the total tax burden as if the taxpayer pays the same tax throughout:

$$\mathbb{CE}_i = \frac{1}{T} \sum_{t=1}^{T} k_{i,t} \bar{x}. \tag{3}$$

The total price can be decomposed into the certainty equivalent and the tax insurance:

$$\mathbb{P}_i = \mathbb{CE}_i + \mathbb{I}_i,$$
$$\frac{1}{T} \sum_{t=1}^{T} k_{i,t} x_{i,t} = \frac{1}{T} \sum_{t=1}^{T} k_{i,t} \bar{x} + \frac{1}{T} \sum_{t=1}^{T} k_{i,t} (x_{i,t} - \bar{x}).$$

Therefore, *Tax insurance* is:

$$\mathbb{I}_i = \mathbb{P}_i - \mathbb{CE}_i = \frac{1}{T} \sum_{t=1}^{T} k_{i,t} (x_{i,t} - \bar{x}). \tag{4}$$

Intuitively, *Tax Insurance* is the weighted average of time-demeaned tax payments (i.e. deviations around the mean), where the relevant weights fluctuate over time depending on the state of the economy. A high payment, i.e., $x_{i,t} - \bar{x}$, gets up-weighted in bad states

Table 1: Tax Insurance Example

	Good Times	Bad Times	Total Price, $\mathbb{P}$	$\mathbb{CE}$	$\mathbb{I}$
Arrow Debreu 1	1	0	\$0.45 (= \$0.9 × 0.5)	\$0.5	-\$0.05
Arrow Debreu 2	0	1	\$0.55 (= \$1.1 × 0.5)	\$0.5	\$0.05
Risk-free Bond	1	1	\$1 (= \$0.45 + \$0.55)	\$1	\$0
Tax payment A	-\$90	-\$110	-\$101 (= −\$90 × 0.45 − \$110 × 0.55)	-\$100	-\$1
Tax payment B	-\$120	-\$80	-\$98 (= −\$120 × 0.45 − \$80 × 0.55)	-\$100	\$2

of the world because having to pay higher taxes during periods of reduced consumption is especially painful. The same high payment gets down-weighted in good states of the world because paying more taxes is less painful when taxpayers earn more and have higher levels of consumption. As a result, a positive $\mathbb{I}$ indicates a lesser risk associated with a given stream of tax payments, and a negative $\mathbb{I}$ indicates a greater risk. A stream of time-invariant payments naturally yields $\mathbb{I} = 0$. Time-invariant weights, corresponding to a taxpayer who does not care about the timing of the payment, also yield $\mathbb{I} = 0$.

3.2 Tax Insurance Example

We summarize the example from the introduction in Table 1. In this example, a good and a bad state of the world occurs with equal probability. The marginal utility of an additional dollar is valued at \$0.9 in the good state of the world and \$1.1 in the bad state of the world. The latter is higher than the former due to diminishing marginal utility. Arrow-Debreu (AD) securities are priced using these marginal utilities and the probability of the state. The AD security that pays \$1 in the good state is worth \$0.45 (= \$0.9 × 0.5) and the one that pays in the bad state is worth \$0.55 (= \$1.1 × 0.5). In this economy, a risk-free bond paying 1 in both states (implicitly combining both AD securities) would be priced at \$1 (= 1 × \$0.45 + 1 × \$0.55). We use these AD securities as the kernel to price the insurance value from income taxes.

The last two rows in table 1 illustrate the insurance value from tax payments in jurisdic-

tions A and B. Tax payments in both jurisdictions have the same expected value, which is also the certainty equivalent, but different timings. In jurisdiction A, individuals pay more taxes in the poor state and pay less in the rich state, which is particularly burdensome. The insurance value is the difference between the total price (-$101) and the certainty equivalent (-$100) and equals -$1. A negative *Tax Insurance* means that individuals would be willing to pay $1 to avoid the risk involved in that tax payment even without changing the expected level of taxation. In contrast, individuals pay more in the rich state and pay less in the poor state, which is relatively less burdensome. Consequently, the *Tax Insurance* in jurisdiction B is $2, where income taxes provide insurance value to individuals.

This stylized example highlights the importance of the timing of the payments. The tax payments in jurisdiction B provide insurance while those in A add risk. This example also highlights the inadequacy of conventional statistics commonly used to measure risk in taxation such as standard deviation. The standard deviation of tax payments in jurisdiction A is 10 while that in B is 20. The tax payments in B are more volatile but reduce the risk whereas those in A are less volatile but increase the risk. Therefore, when assessing the tax risk, one should focus on the timing of the tax liabilities and not just their volatility.

Consider a more detailed example of income tax payments given in Table 2. This example uses the same Arrow-Debreu prices in good times ($0.45) and bad times ($0.55) as the previous example. In the baseline case, the individual earns $45,000 in good economic times and $35,000 in bad times. The income tax is progressive – the income is taxed at 10% up to $40,000 and then at a marginal rate of 25% above $40,000. The average income tax rate is 11.7% in good times and 10% in bad times. The income tax is procyclical – the individual pays more taxes in rich states and pays less in poor states. As a result, this tax regime is less burdensome with a price of -$4,287.50 compared to one that collects the same taxes at all time with a price of -$4,375.00. Consequently, this procyclicality provides insurance valued at $87.50 ($= -\$4,287.50 - (-\$4,375.00)$) by the individual.

We next consider three different scenarios and examine the insurance value provided

Table 2: Income Tax Example of Tax Insurance

	Good Times			Bad Times			Prices		
	Earnings (1)	Tax payment (2)	Ave. tax rate (3)	Earnings (4)	Tax payment (5)	Ave. tax rate (6)	Total price (7)	CE price (8)	Tax Insurance (9)
Baseline	$45,000.00	-$5,250.00	0.117	$35,000.00	-$3,500.00	0.100	-$4,287.50	-$4,375.00	$87.50
Larger shock	$50,000.00	-$6,500.00	0.130	$30,000.00	-$3,000.00	0.100	-$4,575.00	-$4,750.00	$175.00
Higher top rate	$45,000.00	-$5,750.00	0.128	$35,000.00	-$3,500.00	0.100	-$4,512.50	-$4,625.00	$112.50
Higher level	$85,000.00	-$15,250.00	0.179	$75,000.00	-$12,750.00	0.170	-$13,875.00	-$14,000.00	$125.00

by income taxes. First, we consider a larger income shock with which the individual earns $50,000 in good times and $30,000 in bad times. In this case, the average tax rate increases to 13% in good times and remains at 10% in bad times and the expected tax payment increases to $4,750. However, income taxes provide more insurance value, now valued at $175. Second, we consider a more progressive tax system with a higher top rate of 35% instead of 25%. The *Tax Insurance* increases to $112.50. Lastly, we also consider an individual with a higher level of income but the same magnitude of shock in dollar terms. The *Tax Insurance* increases to $125. To summarize, income taxes provide insurance value to individuals, and the value increases when the income is more volatile, the tax schedule is more progressive, and the income level is higher.

3.3 Pricing Kernels

Our measure of *Tax Insurance* is a weighted sum of demeaned tax payments as given in equation 4. These weights constitute a pricing kernel that captures risk preferences. In practice, the pricing kernel can take many different forms. We use the marginal rate of substitution (MRS) in the consumer's Euler equation as the pricing kernel, given as the first term in equation 1. With the assumption of constant relative risk aversion (CRRA) utility,

the pricing kernel is simply

$$k_{i,t} = \left(\frac{c_{i,t-1}}{c_{i,t}} \right)^{\rho}, \tag{5}$$

where ρ is the coefficient of relative risk aversion and $c_{i,t}$ is the consumption of person i at time t. However, it is very difficult to empirically observe individual consumption. Therefore, we use aggregate measures in our empirical analysis. The implicit assumption of using aggregate consumption is that the capital market is complete and all individuals perfectly diversify their consumption risk. As a result, there is no idiosyncratic risk in the cross-section and all individuals equally inherit the aggregate uncertainty in the economy. Although this assumption is unrealistic, it affords us empirical tractability and understates the *Tax Insurance* we calculate. Therefore, *Tax Insurance* calculated with aggregate measures of consumption represents a lower bound of the insurance value for the individual household. We next consider a few choices in the aggregate measures of consumption in Figure 1.

In Panel A of Figure 1, we graph the pricing kernel using U.S. GDP, California GDP, and total consumption as different measures of $c_{i,t}$, with a $\rho = 3.5$, from 1997 to 2021. These measures are highly correlated and weigh payments more in recessionary periods (e.g., 2008) and less in expansionary periods (e.g., 2005). In Panel B, we show how these weights differ using GDP from different states: California (blue), North Carolina (red), and Utah (green). While these series are again highly correlated, some deviations may be important to capture depending on the context in which *Tax insurance* is being calculated.

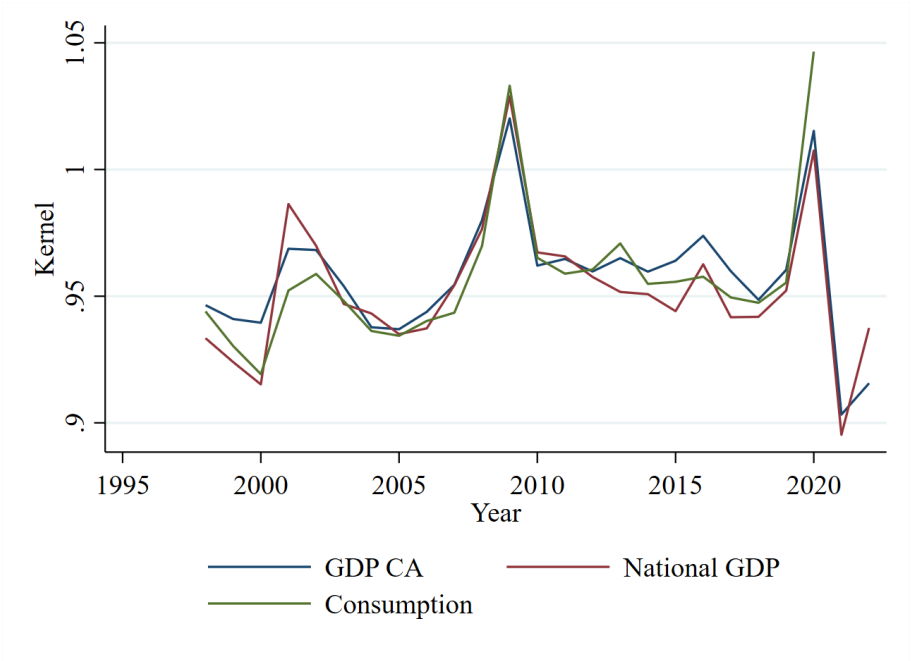
3.4 Tax Insurance Features

We aim for our *Tax insurance* measure to provide a single measure of insurance and risk for future research and reports. Currently, different studies use different measures of insurance and risk making it difficult to draw comparisons across studies. These measures include rolling four-year standard deviations (Poterba, 1994), year-over-year changes (Cornia and Nelson, 2010), and percentage deviations from long-run trends (Seegert, 2015). Boyd (2022)

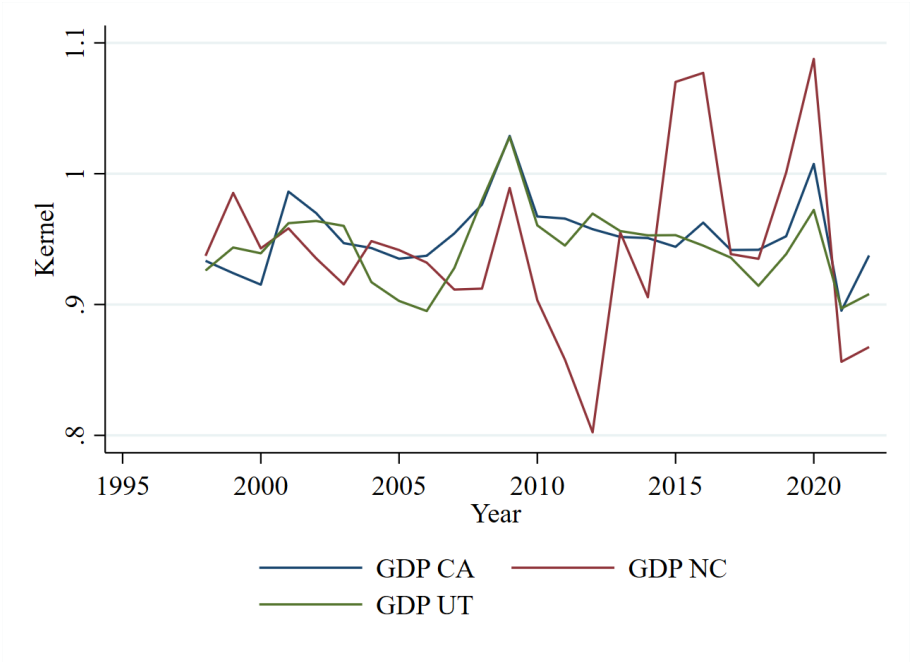
Figure 1: **Pricing Kernel Examples**

This figure provides examples of pricing kernels for calculating the *Tax insurance.*

A) Different Series



B) Different States



offers a discussion of these different measures. Each method presents unique strengths and weaknesses, striking a delicate balance between precision and simplicity.

Our measure of *Tax insurance* captures the strengths of these established approaches while mitigating many of their shortcomings and is on the frontier of simplicity and precision. Notably, the year-over-year changes and rolling standard deviation methods stand out for their inherent intuitiveness and ease of calculation. Their widespread usage can be attributed to these qualities. In contrast, the residual method necessitates researchers to define and often estimate a trend, a task that can be less intuitive, requiring techniques such as the Hodrick-Prescott filter to detrend payments Boyd (2022). Our *Tax insurance* simplifies this process as it is computed as a weighted average, facilitating intuitive and straightforward calculations. Additionally, our measure can be derived from an Arrow-Debreu model, affording it an interpretation akin to a price individuals would be willing to pay to mitigate the associated risks. Our approach, hence, revolves around a standard calculation taught in introductory finance classes and nests standard measures such as skewness.

One noteworthy attribute of the residual method (Seegert, 2015) is its capability to capture the concept of shocks without conflating risk and growth. For instance, consider a linear sequence of tax payments across multiple years. Intuitively, this sequence, being deterministic and evolving smoothly, does not appear to embody substantial risk. Both the residual method and our *Tax insurance* reflect this as having negligible risk. Conversely, risk measurements based on year-over-year changes and rolling standard deviations suggest a significant risk level in such scenarios. This example underscores the tendency of intuitive measures like year-over-year changes and rolling standard deviations to conflate growth and risk, whereas those that do not conflate the two often require complex calculation methods.

Another measure of tax risk with rising popularity is the skewness of tax payments. A large positive skewness indicates that a disproportionately large tax liability is likely to occur, potentially leading to outcomes of distress. We compare skewness to our *Tax Insurance* measure and then demonstrate the strength of our measure in comparison.

Skewness, like our tax risk measure, has a zero-value interpretation. A zero skewness suggests that taxpayers are not more or less likely to have a large or small tax liability. This is because skewness is also a weighted average of demeaned tax payments:

$$\begin{aligned}\mathbb{SK}_i &= \frac{1}{T}\sum_{t=1}^{T}\frac{(x_{i,t}-\bar{x})^3}{\sigma^3} \\ &= \frac{1}{T}\sum_{t=1}^{T}\frac{(x_{i,t}-\bar{x})^2}{\sigma^3}(x_{i,t}-\bar{x}),\end{aligned} \tag{6}$$

where the weighting scheme is $(x_{i,t}-\bar{x})^2/\sigma^3$. When $(x_{i,t}-\bar{x})^2 > \sigma^3$, or the (unsigned) distance from the mean is large, the (signed) deviation gets up-weighted, and vice versa. This weighting scheme suggests that a large deviation is always worse than a small one and therefore the former should be up-weighted and the latter down-weighted. There are no good economic reasons supporting this weighting scheme. Making it worse is the fact that the threshold ($(x_{i,t}-\bar{x})^2 = \sigma^3$) that differentiates large and small distances is completely arbitrary. In stark contrast, our measure re-weights deviations based on the economic state of the world rather than an arbitrary mathematical threshold. Specifically, *Tax insurance* up-weights the deviation when the taxpayer is in the poor state of the world and down-weight it when she is in the rich state of the world. As a result, the weighting scheme of our measure suggests that a large deviation is worse if the individual is in a poor economic state. Therefore, we argue that our measure is grounded in economics and is superior to skewness which is an arbitrary statistic that lacks economic interpretation.

4 Data

We use tax return data from the Individual Public Use Tax File constructed by the Internal Revenue Service (IRS). These data are assembled by the IRS via random sampling of taxpayer identification numbers (TINs) appearing on Forms 1040, 1040A, and 1040EZ from

1979 to 1990. The panel is de-identified, with extensive steps taken to protect taxpayer confidentiality.[1] A newly-assigned unique identifier enables following taxpayers over time. The panel nature of this data is critical for our analysis because it allows us to observe income shocks facing individuals over ten years. We use all of the years produced by the IRS panel. Specifically, while the IRS has more recent data, the panel ends in 1990.

The full sample contains 294,175 returns from 64,219 unique taxpayers. Most taxpayers are in the sample for multiple years (88% more than one year). We restrict our sample to households that appeared in nine or more years (equal to 46% of the data, representing 53 million taxpayers). The descriptive statistics below are for our restricted sample.

The data provides information on wages, adjusted gross income, taxes paid, state, and other variables from the 1040 Forms. For context, the average wage in our sample is $17,675 with a maximum wage of $14,520,000 and a minimum of $0. This wage distribution leads to a tax liability distribution in our sample—including taxpayers with zero taxable income—where the 25th percentile paid $295, the median paid $1,516, and the 75th percentile paid $3,707. The maximum paid was $28,710. The sample contains tax returns from all 50 states and the District of Columbia. Some states have few returns, for example, Alaska has 493 and Wyoming has 477, and we caution placing too much weight on the results for these smaller states. California, for comparison, accounts for 27,350 returns in our sample.

5 Insurance from the Federal Income Tax

5.1 Insurance estimates

We apply our measure of *Tax insurance* to federal income tax payments to estimate insurance values by state. Our estimates suggest that households would have been willing to pay an average of $738 per year for the insurance provided by the individual federal income tax over

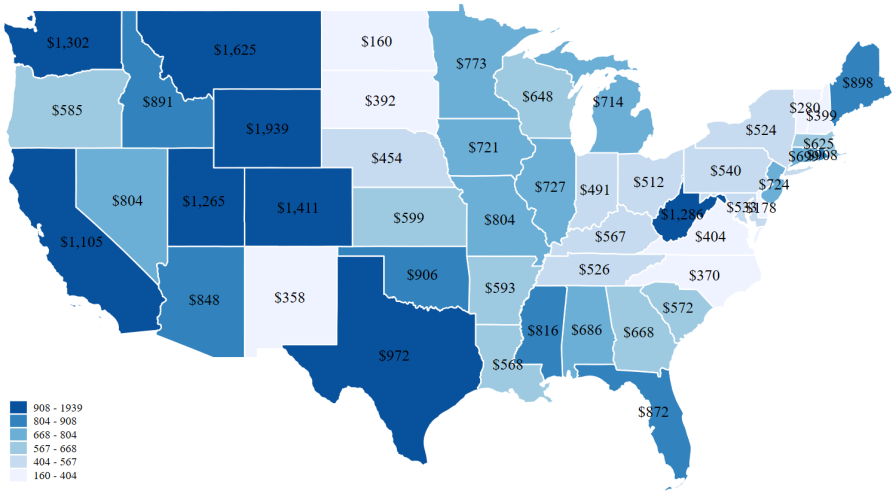
[1] For details on confidentiality procedures, see the IRS documentation at https://www.nber.org/research/data/tax-model-file-documentation.

the period 1979-1990. As shown in Figure 2, however, there is considerable heterogeneity in *Tax insurance* (insurance) across states. The greatest insurance is provided in states in the mountain west, e.g., Wyoming ($1,939), Montana ($1,625), and Colorado ($1,411). The least insurance is provided generally in eastern states, e.g., Vermont ($280) and Delaware ($177), though other states also have low levels of insurance, e.g., North Dakota ($160) and New Mexico ($358). We find similar heterogeneity when considering insurance as a percent of federal income taxes paid (not reported).

Figure 2: **Insurance Federal Income Tax**

This figure uses data from the Statistics of Income on federal income tax returns from 1979 to 1990 to graph *Tax risk* (insurance) by state. We depict states with more insurance and less risk with dark blue and states with less insurance and more risk with light blue.

A) Levels



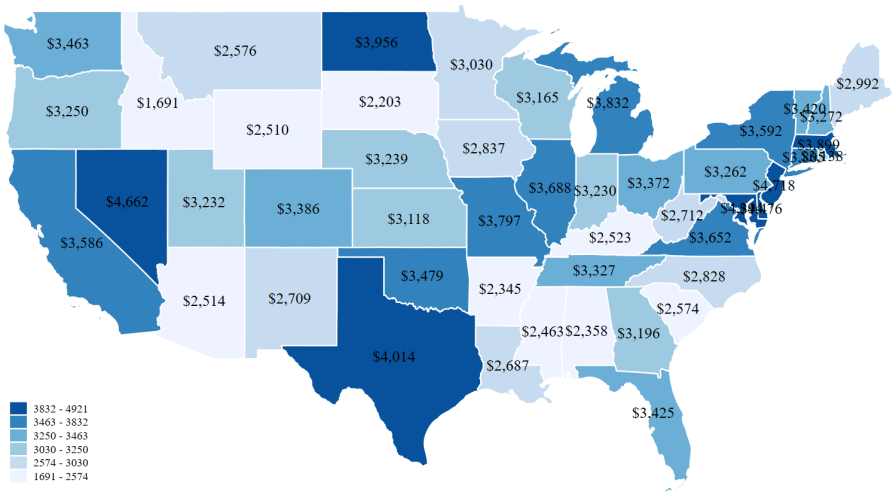
We find that some of the differences across states are due to differences in the variability of taxes paid by state and the average level of taxes paid. In Figure 3, we report the average taxes paid among taxpayers with non-zero liability in Panel A and the average standard deviation of taxes paid in Panel B. Generally, we find a positive relationship between both average taxes paid and the standard deviation thereof with the amount of insurance provided by the federal income tax system. A one standard deviation increase in average taxes paid (roughly $3,463) is associated with an additional $3.40 in insurance. Similarly, we find that a one standard deviation increase in the standard deviation of taxes paid (roughly $1,610) is associated with an additional $17.60 in insurance.

We further explore the relationship between the insurance provided by the federal income tax and the level and variability of wages by simulating different distributions. We begin by simulating 1000 taxpayers' wages across ten years. In the first simulation, we vary the average wage across taxpayers and include year shocks. In the second simulation, we vary

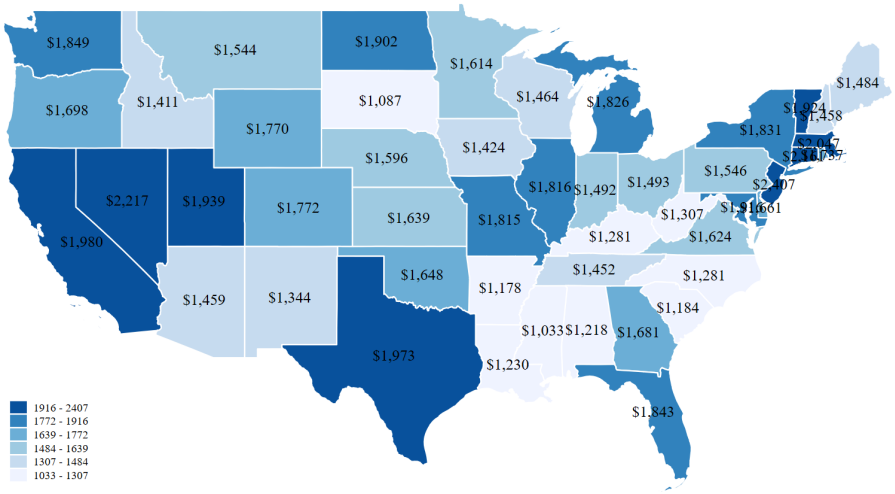
Figure 3: **Tax Variability by State**.

This figure uses data from the Statistics of Income on federal income tax returns from 1979 to 1990 to graph the average taxes paid (Panel A) and the average standard deviation of taxes paid (Panel B). Larger average taxes and standard deviations are depicted in dark blue and smaller average taxes and standard deviations are depicted in light blue.

A) Average Taxes Paid



B) Average Standard Deviation Taxes Paid



the average shock to wages in each year. In both simulations, we use NBER's TAXSIM to calculate the federal income tax. We then use *Tax insurance* to calculate the insurance for each individual. In Figure 4, we report the relationship between the insurance and wages in Panel A and the insurance and standard deviation of wages in Panel B. These simulations support the findings in the data, that higher average wages are associated with greater insurance, and a higher standard deviation of wages is associated with greater insurance.

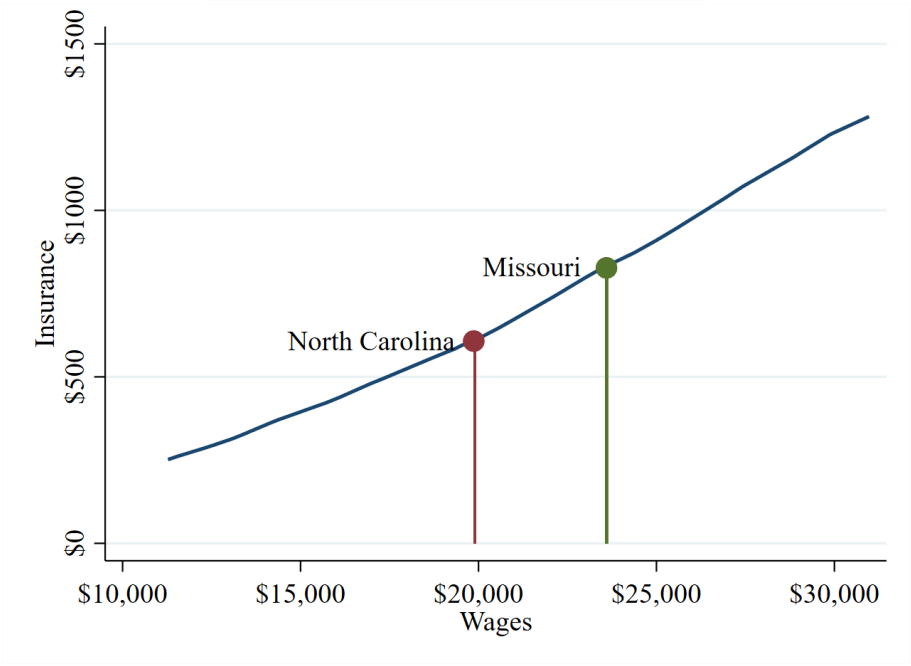
In Figure 4, we show the values of average and standard deviation for North Carolina and Missouri to highlight the effect of these variables on insurance. North Carolina has an average wage of roughly $20,000, and Missouri has an average wage of roughly $24,000. In the simulation, this difference in average wages leads to a difference in the insurance value of $354, holding all else constant. Similarly, North Carolina has an average standard deviation of wages of roughly $8,100 and Missouri of roughly $9,850. This difference in the average standard deviation of wages leads to a difference in insurance of $333, holding all else constant.

This analysis demonstrates that differences in the average level and variability of income can cause states to experience different levels of insurance from the federal income tax. This finding has several policy implications. First, following the large literature on income inequality, this finding suggests that the progressivity of the federal income tax may differentially dampen consumption inequality across states (Dauchy and Balding, 2013; Bargain et al., 2015; Dauchy et al., 2021) Second, following the literature on entrepreneurship and tax-induced risk-taking, this finding suggests that the federal income tax may differentially encourage entrepreneurship and risk-taking across states (Musgrave and Miller, 1948).

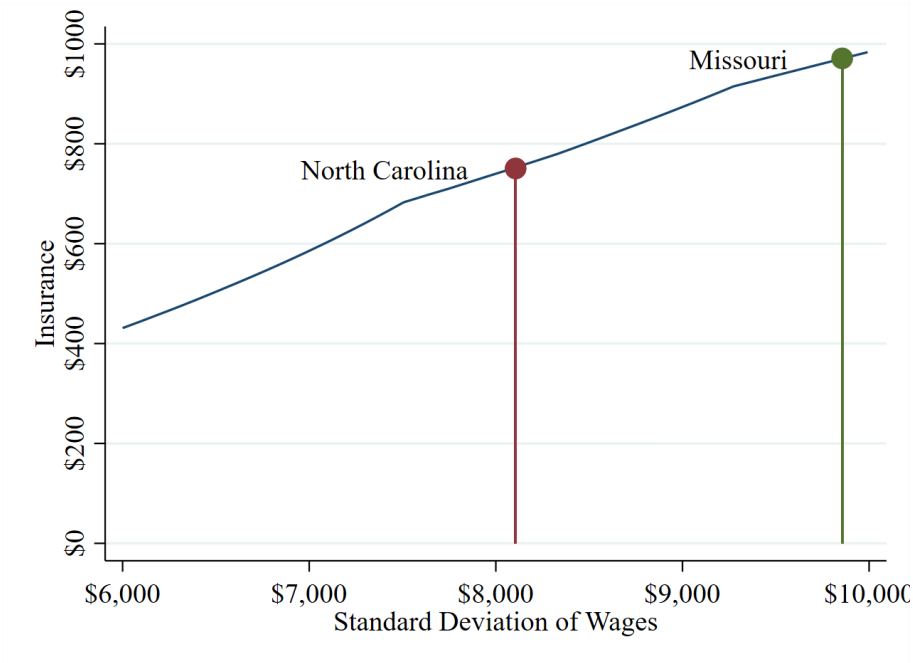
Figure 4: **Risk Simulations**.

This figure depicts *Tax insurance* as a function of wages (Panel A) and the standard deviation of wages (Panel B) that we simulated. The simulation consists of 10,000 taxpayers. We use NBER's TAXSIM program to calculate the federal income tax payments for these taxpayers. For context, we plot North Carolina and Missouri in terms of average wages and standard deviation of wages in the Statistics of Income data from 1979 to 1990.

A) Risk as Average Income Changes



B) Risk as Standard Deviation of Income Changes



6 Insurance from the State Income and Sales Taxes

In this section, we apply our measure of *Tax Insurance* to state income and sales tax. We first describe the data used for the analysis and then examine the value of *Tax Insurance* by state.

6.1 PSID and CEX data

The analysis requires a dataset with income and consumption to consider the impacts of the income and sales tax at the state level. We construct this dataset by combining information from the Consumer Expenditure Survey (CEX) and the Panel Study of Income Dynamics (PSID). We adopt a well-established imputation method, originally introduced by Skinner (1987) and subsequently widely used by researchers, to estimate total consumption within the PSID. This imputation method regresses total consumption on various spending categories, such as food, utilities, and vehicles, common to the PSID and the CEX. The regression coefficients are derived using CEX data. We follow Dauchy et al. (2021) to impute non-durable consumption and select the sample in the cross section. This procedure broadens the coverage in time series and minimizes variation in consumption from non-economic sources such as changes in family structure. The final sample consists of 19,413 observations, a significant reduction from the original 264,810 observations.

We combine the income and consumption data with NBER's TAXSIM program and data on sales tax rates from the Tax Foundation to calculate sales and income tax payments. NBER's TAXSIM program returns tax calculations using federal law from 1960-2023 using data from surveys, such as the PSID.[2] We calculate sales tax payments by interacting the statutory sales tax rate with the portion of consumption imputed from the CEX into the PSID that most closely resembles the sales tax base in the state. For example, Alabama includes grocery food in the sales tax, Wisconsin excludes it, and Utah includes it at a lower rate.

[2]For documentation on NBER's TAXSIM program, see https://www.nber.org/research/data/taxsim.

6.2 Insurance from Income and Sales Tax

We calculate *Tax Insurance* separately for state income taxes and sales taxes. For each household in the sample, we price their state income tax payments and sales tax payments for the entire period of state residency.[3] Specifically, we up-weight deviations from the mean tax payment in bad times and down-weight these deviations in good times. We then take the within-state average to obtain a state-level measure.

Panel (A) of Figure 5 shows the distribution of sales *Tax insurance* across states. We see that sales tax risk is positive across all states.[4] This indicates that the taxes provide insurance. Said differently, residents in these states bear a lower risk associated with their sales tax payments than they would if they had paid a fixed amount (equal to the mean payment) over time. This is because sales taxes are inherently pro-cyclical—people tend to purchase more goods and services (and therefore pay more sales taxes) in good times, and vice versa when economic conditions are bad. For example, taxpayers in Illinois would be willing to pay over $350 per year to maintain the pro-cyclicality of their sales tax payments compared to paying a constant amount equal to their expected annual payment regardless of their income and consumption.

A similar pattern appears in the distribution of state income tax risk, as shown in panel (B) of Figure 5. Most states' income tax systems also produce positive risk prices; however, state rankings of *Tax insurance* differ across tax instruments.[5] Illinois thus provides the greatest average insurance against sales tax payments across households in the PSID, yet the amount of insurance afforded by its income tax regime is near the median amount.

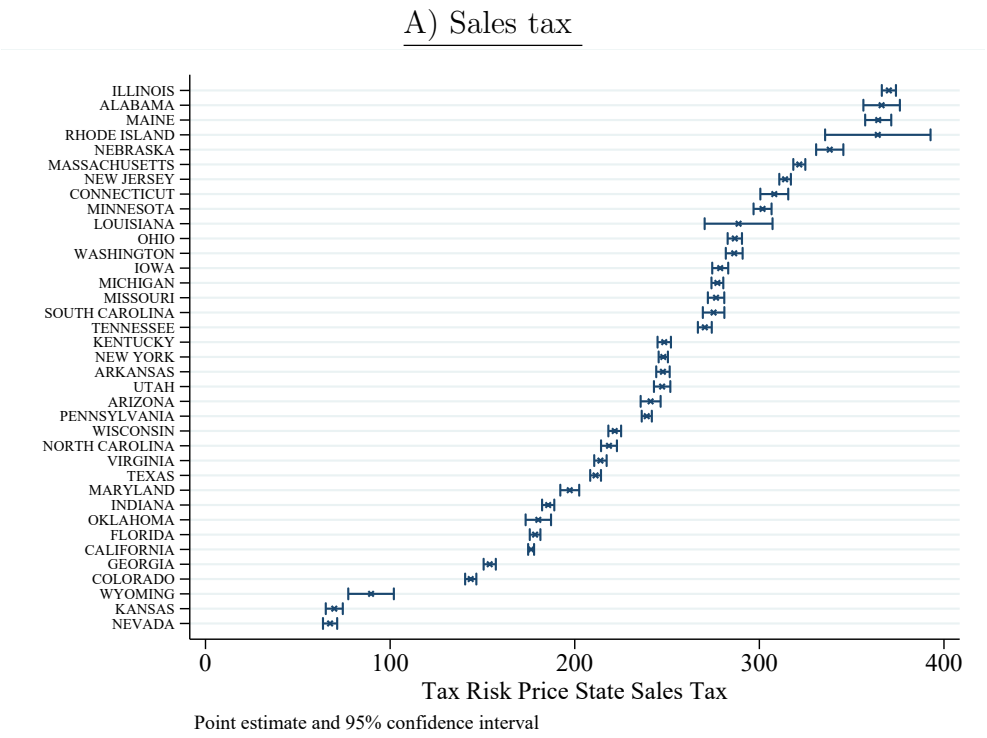
[3] Survey respondents that move between states are treated as independent units.

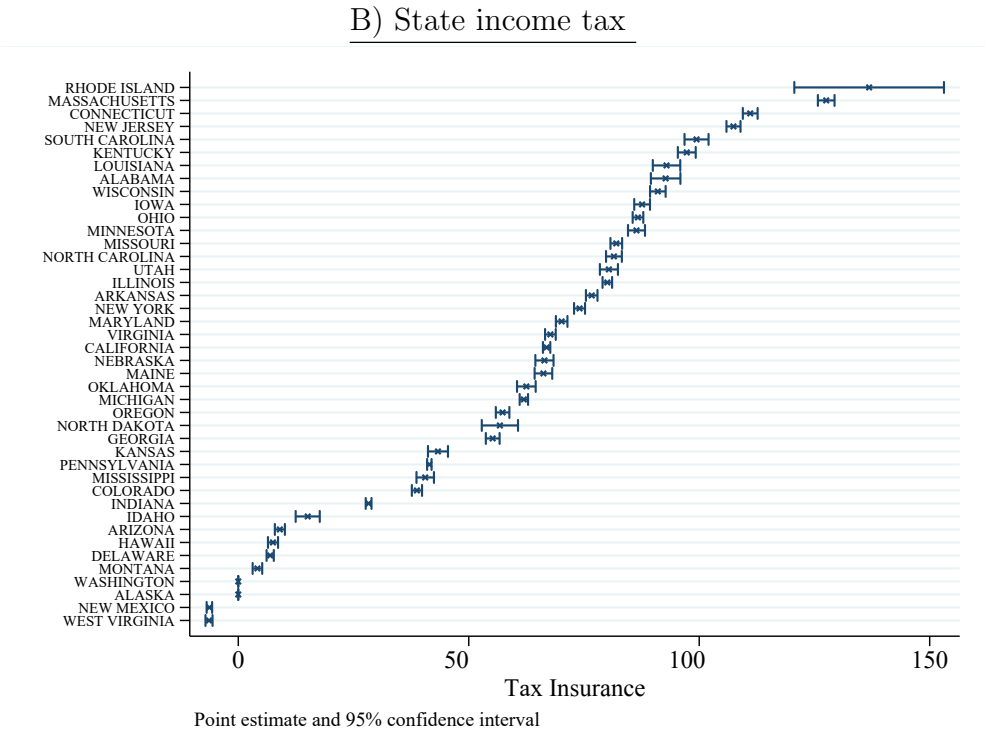
[4] Several states, including Alaska, Delaware, Montana, New Hampshire, and Oregon do not impose any sales taxes, hence *Tax Risk* is undefined for these states. We also exclude the District of Columbia, Hawaii, Idaho, Mississippi, New Mexico, North Dakota, South Dakota, Vermont, and West Virginia from panel (A) of Figure 5 due to these states having small numbers (i.e., single digits) of survey respondents in the PSID for which we can calculate sales tax liabilities, such that average *Tax insurance* among these respondents may not necessarily be representative of the state population.

[5] Alaska, Florida, Nevada, South Dakota, Tennessee, Texas, and Wyoming do not have state income taxes. We also exclude the District of Columbia, New Hampshire, and Vermont from panel (B) of Figure 5 due to these states having small numbers (i.e., single digits) of survey respondents in the PSID for which we can calculate income tax liabilities (and hence, representative measures of *Tax insurance*).

Figure 5: **Income Tax Insurance by State**

This figure calculates the *Tax Insurance* for income and sales and tax payments of individuals in the Panel Study of Income Dynamics (PSID) with imputed consumption using the Consumer Expenditure Survey (CEX). Panel A depicts the average *Tax Insurance* of sales tax payments by state. Panel B depicts the average *Tax Insurance* of income tax payments by state.





Overall, *Tax insurance* at the household level from sales taxes is also generally higher (i.e., providing greater insurance value) than for state income taxes. For both types of state taxes, *Tax Insurance*—the amount of tax cut that people would demand in exchange for facing a flat schedule of tax payments every year—is substantial.

7 Conclusion

This study addresses a crucial yet often understudied aspect of fiscal policy: the measurement and pricing of tax insurance and risk. Despite the recognized significance of quantifying the progressivity of a tax system or the risk associated with a tax system, there has been a lack of consensus on a precise definition of tax insurance and risk. We fill this gap by introducing a novel method for pricing tax insurance that encapsulates key risk components while remaining straightforward to implement.

The core of our method lies in treating tax payments as analogous to returns from a financial asset and applying the same pricing principles used in asset pricing models like the Arrow-Debreu model. By viewing tax payments through this lens, we can calculate a measure of *Tax insurance*, which quantifies the covariance between tax payments and the value of a dollar in different states of the world. This approach goes beyond conventional measures of tax risk or uncertainty, such as the variance of tax payments, by incorporating elements crucial to understanding the real economic effects of tax insurance and risk, including its asymmetry and its correlation with different states of the world.

This measure of tax insurance has several desirable characteristics: it is easy to implement, provides economic content by assigning a price to the risk involved, and can be applied to a wide range of contexts. We demonstrate the applicability of *Tax insurance* in by evaluating the insurance provided to households by the federal income tax. *Tax insurance* sheds light on the scale and distribution of tax insurance and its implications for decision-making.

The findings presented in this study have significant policy implications. The results

show that households face substantial risk and the federal income tax absorbs some this risk. Understanding and quantifying this risk can aid in making more informed fiscal decisions. Tax insurance is not uniform across jurisdictions, and its impact varies based on the specifics of the tax system and the economic conditions within each state. Moreover, the analysis reveals that the federal income tax provides differing levels of insurance to households across states, raising intriguing questions about the broader economic effects of this variation. We intend for our newly-developed measure of *Tax insurance* and our dissemination of Stata code to facilitate its calculation to spawn wide-ranging inquiries into quantifying these economic effects.

The insights provided by this study contribute to a more comprehensive understanding of the multifaceted nature of tax risk and its implications for economic agents and government entities alike.

References

Abowd, J. M. and D. Card (1989, March). On the covariance structure of earnings and hours changes. *Econometrica 57*(2), 411–45.

Aguiar, M. and M. Bils (2015). Has consumption inequality mirrored income inequality? *American Economic Review 105*(9), 2725–56.

Altonji, J. G. and A. Siow (1987). Testing the response of consumption to income changes with (noisy) panel data. *The Quarterly Journal of Economics 102*(2), 293–328.

Armour, P., R. V. Burkhauser, and J. Larrimore (2016). Using the pareto distribution to improve estimates of top coded earnings. *Economic Inquiry 54*(2), 1263–1273.

Attanasio, O. P., E. Hurst, and L. Pistaferri (2014). The evolution of income, consumption, and leisure inequality in the united states, 1980–2010. In *Improving the Measurement of Consumer Expenditures*, pp. 100–140. National Bureau of Economic Research, Inc.

Attanasio, O. P. and N. Pavoni (2011, 07). Risk sharing in private information models with asset accumulation: Explaining the excess smoothness of consumption. *Econometrica 79*(4), 1027–1068.

Attanasio, O. P. and G. Weber (1995, December). Is consumption growth consistent with intertemporal optimization? evidence from the consumer expenditure survey. *Journal of Political Economy 103*(6), 1121–57.

Auerbach, A. J. and D. Feenberg (2000). The significance of federal taxes as automatic stabilizers. *Journal of Economic Perspectives 14*(3), 37–56.

Autor, D. H., L. F. Katz, and M. S. Kearney (2008). Trends in us wage inequality: Revising the revisionists. *The Review of Economics and Statistics 90*(2), 300–323.

Bargain, O., M. Dolls, H. Immervoll, D. Neumann, A. Peichl, N. Pestel, and S. Siegloch (2015). Tax policy and income inequality in the united states, 1979–2007. *Economic Inquiry 53*(2), 1061–1085.

Bloom, N., S. Bond, and J. Van Reenen (2007). Uncertainty and investment dynamics. *Review of Economic Studies 74*(2), 391–415.

Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *American Economic Review 98*(5), 1887–1921.

Boyd, D. (2022). State tax revenue volatility and its impact on state governments. The Pew Charitable Trust white paper.

Bradley, S., D. Huang, and N. Seegert (2023). Asymmetric risk of housing distress from property tax limitations. Working paper.

Carroll, C. D. and M. S. Kimball (2001, October). Liquidity constraints and precautionary saving. *NBER Working Paper* (8496).

Chetty, R. (2006). A bound on risk aversion using labor supply elasticities. *NBER Working Paper* (w12067).

Cornia, G. C. and R. D. Nelson (2010). State tax revenue growth and volatility. *Federal Reserve Bank of St. Louis Regional Economic Development 6*(1), 23–58.

Cutler, D. M. and L. F. Katz (1992, May). Rising inequality? changes in the distribution of income and consumption in the 1980's. *American Economic Review 82*(2), 546–51.

Dauchy, E., F. Navarro-Sanchez, and N. Seegert (2021). Taxation and inequality: Active and passive channels. *Review of Economic Dynamics 42*, 156–177.

Dauchy, E. P. and C. Balding (2013). Federal income tax revenue volatility since 1966. *Working Paper*.

Debacker, J., B. Heim, V. Panousi, S. Ramnath, and I. Vidangos (2013). Rising inequality: Transitory or persistent? new evidence from a panel of us tax returns. *Brookings Papers on Economic Activity*, 67–122.

Domar, E. D. and R. A. Musgrave (1944). Proportional income taxation and risk-taking. *The Quarterly Journal of Economics 58*(3), 388–422.

Graham, J. R., M. Hanlon, T. Shevlin, and N. Shroff (2014). Incentives for tax planning and avoidance: Evidence from the field. *The Accounting Review 89*(3), 991–1023.

Hassett, K. A. and G. E. Metcalf (1999). Investment with uncertain tax policy: Does random tax policy discourage investment. *The Economic Journal 109*(457), 372–393.

Heathcote, J., F. Perri, and G. L. Violante (2010). Unequal we stand: An empirical analysis of economic inequality in the united states, 1967–2006. *Review of Economic Dynamics 13*(1), 15–51.

Heathcote, J., F. Perri, G. L. Violante, and L. Zhang (2023). More unequal we stand? inequality dynamics in the united states, 1967–2021. *Review of Economic Dynamics 50*, 235–266.

Jacob, M., K. Wentland, and S. A. Wentland (2022). Real effects of tax uncertainty: Evidence from firm capital investments. *Management Science 68*(6), 4065–4089.

Kaplan, G. and G. L. Violante (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics 2*(4), 53–87.

Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239.

Larrimore, J., J. Mortenson, and D. Splinter (2015). Income and earnings mobility in us tax data.

MaCurdy, T. E. (1982, January). The use of time series processes to model the error structure of earnings in a longitudinal data analysis. *Journal of Econometrics 18*(1), 83–114.

Meghir, C. and L. Pistaferri (2004, 01). Income variance dynamics and heterogeneity. *Econometrica 72*(1), 1–32.

Musgrave, R. A. and M. H. Miller (1948, March). Built-in flexibility. *American Economic Review 38*(1), 122–28.

Pechman, J. A. (1973). Responsiveness of the federal individual income tax to changes in income. *Brookings Papers on Economic Activity 4*(2), 385–428.

Poterba, J. M. (1994). State responses to fiscal crises: The effects of budgetary institutions and politics. *Journal of Political Economy*, 799–821.

Seegert, N. (2015). The performance of state tax portfolios during and after the great recession. *National Tax Journal 68*(4), 901–918.

Skinner, J. (1987). A superior measure of consumption from the panel study of income dynamics. *Economic Letters 23*(2), 213–16.

Splinter, D. (2012). *Income Variability: Effects on US Income Inequality and Tax Progressivity*. Rice University.

Zeldes, S. P. (1989). Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy 97*(2), 305–46.