

Implicit Taxes and Measured Progressivity

Claude

Experimental draft — not for citation or circulation · August 7, 2026

Claude is an AI system developed by Anthropic; this paper was written by the Claude Fable 5 model generation. This paper was commissioned by Jason Furman.

ABSTRACT

Municipal bondholders bear a tax no distribution table records: the yield they forgo because the exemption is capitalized into prices. Every maintained United States distributional framework counts the interest as income and records zero tax against it. We measure the implicit tax at the household level, combining Survey of Consumer Finances portfolios, marginal tax rates computed on the interest margin with TAXSIM, and a market-implied tax rate identified from the short end of the municipal term structure (38.0 percent, following Longstaff 2011). Correcting average effective federal income tax rates by percentile raises the 99th-to-99.9th percentile's rate by 0.34 percentage points at the anchor — 0.09 to 0.39 across the paper's grid of market-implied rates — with essentially no change below the 90th percentile, and adds \$34 to \$50 billion of unrecorded burden for tax year 2021. The survey's aggregate tax-exempt interest exceeds the administrative record by a factor of 1.47, and the correction is close to linear in that aggregate, so the same percentile correction computed on the IRS base is 0.23 points rather than 0.34; the distributional shape is unchanged. Decomposing each holder's statutory rate into market-priced burden and inframarginal rent, we find the short-end market prices the average municipal bond dollar closely: aggregate rent in 2021 was within a billion dollars of zero under the two block-average constructions that pair holder and market rates on the same tax basis, and ran from $-\$8.8$ to $+\$12.4$ billion across the full set of consistent pairings — small in every case against the burden. Rent is therefore not what the short-end anchor finds; it appears only under the lower rates that sit further out the curve, which this paper carries as stated scenarios rather than as readings of what 2021 holders were paid by maturity. Two neighboring corrections are sized here but not estimated: an owner-occupied housing correction would run the other way, landing in the middle of

the distribution rather than at the top, and a corporate-equity correction cannot currently be signed at all. Neither is a bound, and neither is computed by percentile.

Every number in this paper regenerates from public data by a single command; see the reproducibility statement in Section 4 and Appendix A5.

1. Introduction

Consider a household in the top 0.1 percent of the income distribution in 2021. In the Survey of Consumer Finances, such a household more likely than not reports tax-exempt interest income: 55.5 percent of tax units in that group do, against 2.4 percent of all tax units.¹ The household's tax return reports that interest on line 2a of Form 1040, where it is counted as income and taxed at zero. Every maintained distributional framework in the United States records the same thing: municipal bond interest enters the income base at the yield the household received, and no tax is recorded against it. The measured conclusion is that this income escapes tax.

The household bears a real burden anyway. When it bought the bond, it accepted a yield below what a comparable taxable bond paid. That forgone yield is a tax in everything but remittance: the household gives up income because of the tax system, and the issuing government collects the difference as a lower borrowing cost. Galper and Toder (1984) named it four decades ago: savers bid down the yield of tax-exempt assets, and "this decline in the before-tax yield constitutes an implicit tax." What the holder remits to the Treasury is zero; what she pays is not.

No agency corrects for this, and the most direct confirmation comes from authors who defend the practice. Saez and Zucman (2024), in a methodological paper that treats municipal bonds as a litmus test, write that a tax-incidence approach "should logically assign muni investors a higher counterfactual pre-tax income... and then fictitious taxes they do not pay," and that "to the best of our knowledge... in practice US agencies and think-tanks follow our methodology and do not gross up muni interest income by fictitious taxes." They regard the convention as the right default for tax-reform analysis; Section 2 takes up their argument directly. Here it serves as certification of the gap: the correction this paper performs is recognized as the logical incidence treatment, and nobody does it.

This paper does it, for the asset where it can be done from market prices. We combine household portfolios and incomes from the 2022 Survey of Consumer Finances (income year 2021), statutory marginal tax rates computed for each tax unit with NBER's TAXSIM calculator, and a market-implied municipal tax rate identified from the term structure of municipal yields. With these three inputs we recompute average effective federal income tax rates by income

¹ The top-0.1 percent cell rests on a few hundred oversampled observations, so the participation rate is a point estimate with nontrivial sampling error; the SCF also excludes the Forbes 400 by design, so the very top of the group is absent. The concentration pattern across groups, which is what the argument uses, is robust. Section 4 describes the data.

percentile, with and without the implicit tax on municipal bonds. Throughout, the statutory rate decomposes as $\tau_i = \tau^* + (\tau_i - \tau^*)$: the market-implied rate τ^* is the burden every holder pays through the depressed yield, and the excess of the holder's own rate over the market rate is an inframarginal rent, a windfall the tables also fail to record. We report both parts separately, because they answer different questions: the first corrects measured burdens, the second measures what the exemption gives away.

The paper reports three findings. First, the correction is concentrated exactly where measured top rates bind, and it is material there. Across the paper's grid of market-implied rates — the term structure plus the working-paper version of the anchor — the correction to the 99th-to-99.9th percentile group, the paper's headline group (the top tenth is reported separately), runs from 0.09 to 0.39 percentage points; at the anchor — the short-end identified rate of 38.0 percent (Longstaff 2011, the published estimate; the working-paper version reported 41.6) — it rises by 0.34 points, and the largest correction belongs to the 95th to 99th group, at 0.50. Below the 90th percentile no group moves by more than a tenth of a point. The anchor is near the top of the paper's own grid, chosen for identification cleanliness rather than size; its federal component gives 0.30 points.

Second, the aggregate unrecorded burden is material under every construction the paper prefers, and never less than \$14 billion on the survey income base. Grossing up reported tax-exempt interest at the anchor gives \$50 billion for tax year 2021 on the survey income base, \$34 billion scaled to the administrative record, and \$44 billion at the federal-component rate; each construction is stated with its base in Section 5. At today's stocks and yields the range is \$47 billion (the observed ten-year spread) to \$58 billion (the anchor crossed with the ten-year yield, an upper-bound pairing). For scale, Treasury scores the entire tax expenditure for public-purpose municipal interest at \$24 billion for fiscal 2024, and calendar-2021 federal personal current taxes were \$2,128 billion, so the unrecorded burden is roughly 2.3 percent of federal personal current taxes that year.

Third, whether the exemption is a windfall to the rich turns out to be a question about maturity — one no distribution table has recorded in either direction. Measured against the short-end anchor, with holder and market rates paired on the same tax basis and holders' rates computed as average rates on the full counterfactual taxable block, the aggregate inframarginal rent in 2021 was approximately zero: plus \$1.0 billion on the expected-state pairing, minus \$0.9

billion on the federal-only pairing, against \$50 billion of burden. Pairings that compute holders' rates at the margin rather than on the block widen the range to minus \$8.8 through plus \$12.4 billion, which is still small against the burden and still not the large windfall the exemption is usually said to deliver. The short-end market prices the average municipal bond dollar closely: no aggregate windfall. At the thirty-year implied rate, by contrast, most of the statutory rate is rent.

The measurement stakes are specific. After-tax income is recorded correctly under either convention, because the household's cash income and its remitted tax are both observed. The error is confined to the ratio of tax to income, the effective rate, and effective rates by income group are precisely the object on which progressivity claims are made. On the municipal margin, correcting the ratio makes the tax system look more progressive than the tables say. We state that plainly, and we also state its converse: correcting only municipal bonds, the asset held at the top, while leaving other capitalized preferences uncorrected, biases the overall picture toward more progressivity. Section 6 sizes the two large neighbors, owner-occupied housing and corporate equity, and states which way each pushes, rather than pretending the correction is complete.

The contribution is a corrected effective-rate schedule built from market prices, with the burden and the rent reported separately by percentile. Implicit taxes have met progressivity measurement before. Iglarsh and Allan (2003) imputed a municipal implicit tax from tax-return data and recomputed aggregate progressivity indices; we know that paper only from its published abstract, and the differences we assert — a market-identified rate rather than an assumed spread, a percentile-level rate schedule rather than an index, and the burden-rent decomposition — are conditional on a full read, which is the first item for the coauthor round.² Galper, Rosenberg, Rueben, and Toder (2013) distributed the benefit of the exemption under incidence assumptions, the other side of the ledger from measured burdens. Neither produces what the distribution tables would need: corrected rates by income group.

The data marriage is not new either, which is a point in its favor. Poterba and Ramírez Verdugo (2011) combined SCF portfolios with TAXSIM marginal rates to estimate the revenue cost of the exemption under portfolio-substitution counterfactuals. We use the same marriage for a different object: measured burden rather than counterfactual revenue. The founding empirical fact goes back further. Feenberg and Poterba (1991), using the first tax returns on which

² Advances in Taxation 15 (2003); paywalled at the time of writing. The engagement here is deliberately limited to what the abstract supports.

tax-exempt interest had to be reported, found the holdings concentrated at high marginal rates. Municipal bonds sit at the top of the distribution; that is why recording their tax as zero moves measured progressivity where it is most watched.

Everything turns on the implicit tax rate, and the naive measure of it fails. The ratio of municipal to Treasury yields embeds credit risk, liquidity, and call provisions along with tax; read off the Bond Buyer index, it implies a marginal investor whose tax bracket turned negative after 2008. Longstaff (2011) provides the identification the paper leans on: at the very short end of the municipal swap market, where credit and call effects are negligible and liquidity is stripped through the swap term structure, a naive reading in the mid-teens becomes an identified marginal tax rate of 38.0 percent. We anchor the headline on that rate and treat the maturity term structure — down to roughly 14 percent at thirty years today — as the robustness dimension, because the long end of the municipal market is a genuinely open puzzle, not a nuisance to be assumed away. Section 3 takes the puzzle up directly.

Three limits are stated here and maintained throughout. The exercise is partial-equilibrium: it reallocates measured burden at observed prices and makes no claim about what yields would be in a different tax system; Section 7 takes up the general-equilibrium question, including the transfer to tax-exempt institutions that a Miller-style gross-up would imply and that no one has quantified — the open-economy stance implies no such economy-wide transfer, and the paper declines to choose between them. The income year is 2021, when pandemic-era credits pushed measured rates at the bottom of the distribution to unusual lows; the municipal correction barely interacts with that, but the baseline schedule wears it. And the two available measures of aggregate tax-exempt interest disagree: the survey base is \$81.4 billion, the administrative base \$55.5 billion. We carry both rather than adjudicate, and because the correction is close to linear in tax-exempt interest, the choice moves the rate correction roughly in proportion — 0.34 points on the survey base against 0.23 on the administrative one, with the shape of the schedule unchanged.

Section 2 sets out the accounting identity and documents the convention it corrects. Section 3 identifies the implicit tax rate. Section 4 describes the data and the construction. Section 5 reports the corrected schedules, the aggregates, and the rent decomposition. Section 6 sizes the housing and equity corrections and states the direction of each. Section 7 states what the correction assumes about incidence and what remains unknown. Section 8 concludes.

2. The remittance convention and the decomposition

Every maintained distributional framework in the United States measures tax burdens by taxes remitted. Remittance and burden are different objects — the point is old and general (Slemrod 2008) (verify) — but for most income the difference is harmless in the respect that matters here: the wage earner's remitted liability is levied on income the tax system did not first depress, so recording it captures the direct burden even though incidence may shift it to others. (Withholding is a third object again — a provisional payment reconciled at filing — and it is the annual liability, not the withholding, that these frameworks and this paper use.) For income from an asset whose tax preference is capitalized into its price, remittance and burden come apart at the source: the price itself has moved. The preference changes the pre-tax income the asset generates, so a table that starts from observed income and remitted tax records a burden of zero on a holder who is in fact paying every year, in forgone yield, for the privilege of the exemption.

Table 1 documents the convention across the six frameworks that produce the distribution tables policymakers use. The Congressional Budget Office includes tax-exempt interest in its capital income concept and records no tax against it (Congressional Budget Office 2024). Treasury's Office of Tax Analysis includes it in cash income on the same terms (Cronin 2022). The Joint Committee on Taxation includes it in expanded income (Joint Committee on Taxation 2013) (verify). Auten and Splinter (2024) include it in pre-tax income, as do Piketty, Saez, and Zucman (2018) in distributional national accounts; Splinter (2020) carries the same convention into progressivity indices, and the Tax Policy Center's microsimulation does the same in expanded cash income. None makes a capitalization adjustment. The uniformity is the finding: this is a shared convention, not an oversight by one agency.

Table 1. How the maintained distributional frameworks treat tax-exempt interest. Sources: the methodology documents cited in the text, reviewed for this paper.

Framework	Tax-exempt interest in income?	Tax recorded against it	Capitalization adjustment
CBO, Distribution of Household Income	Yes (capital income, App. D)	None	None
Treasury OTA (TP-8, Cronin 2022)	Yes (cash income)	None	None
JCT distributional analyses	Yes (expanded income)	None	None
Auten–Splinter (2024)	Yes (pre-tax income)	None	None
Piketty–Saez–Zucman DINA	Yes (national income)	None	None
Tax Policy Center microsimulation	Yes (expanded cash income)	None	None

The correction rests on one definition and one identity. For municipal and taxable bonds of comparable maturity and risk, define the market-implied tax rate $\tau^* = 1 - r_m/r_t$, where r_m and r_t are the tax-exempt and taxable yields. A holder whose statutory marginal rate on interest is τ_i and who could hold either bond earns the same after-tax return from the municipal bond as from a taxable bond taxed at rate τ^* . Against a risk- and maturity-matched taxable twin, a voluntary holder must satisfy $\tau_i \geq \tau^*$, and her statutory rate splits into two parts:

$$\tau_i = \tau^* + (\tau_i - \tau^*).$$

The first part is a burden: every holder gives up τ^* of the taxable yield, whoever she is. The second part is an inframarginal rent: a holder whose own rate exceeds the market-clearing rate keeps the difference, a windfall priced by some other investor's lower bracket. The inequality's conditions matter for what follows. Real holders do not face matched twins: their bonds differ from Treasuries in maturity and risk, their tax rates include state components a federal calculation misses, and some hold munis for reasons other than tax. Measured violations of $\tau_i \geq \tau^*$ are therefore diagnostic rather than paradoxical, and Section 5 reads them that way. One further objection belongs here: a holder whose rate sits below the market's could avoid the wedge by selling, so is an avoidable wedge a tax borne? It differs from an ordinary dominated portfolio choice in the respect that matters for burden accounting: the wedge is created by the tax law, and it is collected by a government — the issuer's lower borrowing cost is revenue, whoever remits it. Section 7 returns to who collects.

The operational form works from observed income. A holder who reports tax-exempt interest y would have needed pre-tax income $y/(1 - \tau^*)$ to net the same amount from a taxable bond taxed at the market rate. The implicit tax is the difference, $y \cdot \tau^*/(1 - \tau^*)$, and the corrected effective rate adds it to both sides of the ratio:

$$\text{ETR corrected} = (T + y \cdot \tau^*/(1 - \tau^*)) / (Y + y \cdot \tau^*/(1 - \tau^*)),$$

where T is remitted federal income tax and Y is the household's expanded income with tax-exempt interest counted at the observed yield. After-tax income is unchanged by the correction: the household keeps $Y - T$ under either convention, because the implicit tax added to the numerator is exactly the pre-tax income added to the denominator: the entire measurement

error sits in the rate, not the level.³

A worked example fixes magnitudes, with the rates it pairs stated each time. A household reporting \$100,000 of municipal interest, corrected at the federal-component rate of 35 percent, has a pre-tax-equivalent income of \$153,846 and an implicit tax of \$53,846; at the 2021 federal top rate on interest of 40.8 percent, its rent is \$8,923. The same income at today's thirty-year implied rate of 14.5 percent carries an implicit tax of \$16,959 and a rent of \$30,760. And against the full combined anchor of 38.0 percent, a holder whose combined rate matches the dollar-weighted average of 32.4 percent pays more through the yield than the statute would have charged: rent of minus \$9,032 — the diagnostic case Section 5 measures in the population, where it turns on how the counterfactual rate is computed. The same contrast runs through everything that follows: at the short end the exemption is all burden and no windfall; at the long end it is mostly windfall.

What counts as burden needs one more sentence, because the municipal market since 2008 makes it easy to get wrong. The burden is the tax component of the yield give-up. In stretches where municipal bonds out-yield Treasuries, the extra yield compensates credit and liquidity risk that the holder actually bears; compensation for risk is not taxable-equivalent income, and a negative naive spread is not a negative burden. It is contamination, and stripping it out is precisely what the identification in Section 3 is for.

Rates are what progressivity is measured on: average-rate schedules directly, and any index or top-share claim built on such a schedule inherits the convention with it.

The convention errs in both directions at once, which is why the decomposition matters. Recording zero tax understates the burden at the top: the system is more progressive on the municipal margin than the tables say. Recording nothing about the rent understates what the exemption gives away to holders whose brackets exceed the market rate. A paper that reported only the first would read as a brief for the proposition that the rich pay more than you think; one that reported only the second, as a brief against the exemption. Reporting both, by percentile, is measurement. Which part dominates is an empirical question about τ^* , and Section 5 shows the answer depends on maturity.

The idea has a lineage that makes its absence from the tables harder to explain. Galper and Toder (1984) built a "full effective tax rate" against the return available absent the preference

³ The algebra is in Appendix A1.

and computed illustrative households. Fullerton (1984) cataloged the competing effective-tax-rate concepts and insisted the choice follow the question being asked. Scholes and Wolfson (1992) systematized implicit taxes at the level of firms and asset classes, and Williams (2002) models the multi-clientele equilibrium that keeps the market-clearing implicit rate below the top statutory rate — the theoretical face of this paper's rent term. What never happened, on the evidence of the methodology documents in Table 1 and the searches behind this paper, is the carry-through to an economy-wide distribution table.

Saez and Zucman (2024) give the convention its best defense, and the disagreement is narrower than it looks. Their argument: distributional analysis should allocate taxes actually paid at observed incomes, because counterfactual gross-ups import incidence assumptions and blur what a reform would collect. For scoring reforms, we agree. But progressivity statements are burden statements about current law, and their own municipal litmus test concedes the point on which this paper rests: current-tax analysis, they write, "captures that the pre-tax incomes of the owners of munis are depressed, although it does not single out taxes as the culprit." Measurement of who bears burdens needs the culprit. The implicit tax is a burden the current law imposes and current prices reveal; assigning it requires no counterfactual behavior, only the market's own rate.

The correction accordingly needs three observables and nothing else: holdings or income from the asset, statutory rates household by household, and τ^* . The first two are routine. The third is where casual treatments fail, and the next section is about getting it right.

One scope note belongs before the identification. The schedule this paper corrects is the federal individual income tax schedule. The market rate τ^* that the correction applies, however, is what the market charges, and Section 3 argues that rate is best read as embedding state as well as federal tax. The two constructions that follow from that tension — correcting the federal schedule with the full market rate, or with its federal component only — are both computed, both reported, and named wherever they appear.

3. Identifying the implicit tax rate

The natural first move is to read τ^* off index yields, and it fails in an instructive way. Computed from the Bond Buyer 20-Bond index over the 20-year Treasury, the naive implied rate averaged 10.9 percent in the 1950s, 18.5 percent in the 1960s, 19.6 percent in the 1970s, 15.7 percent in the 1980s, and 14.6 percent in the 1990s. In the 2000s it fell to 3.6 percent, and in the 2010s it averaged -35.1 percent. A marginal investor in a negative tax bracket does not exist. The ratio embeds credit risk, liquidity, and the call option along with tax, and after 2008 those components swamped the tax component entirely.

The full series, plotted against the top statutory rate on interest income in Section 5's historical exhibit, makes the pathology visible: long stretches after 2008 have municipal bonds out-yielding Treasuries — dealer balance sheets shrank, taxable Build America Bonds changed the supply mix, and in March 2020 the municipal market dislocated outright, with ratios far above one. Whatever a reader makes of any particular episode, the lesson is the one Section 2 stated: the naive spread is not a burden measure, and a correction built on it would produce nonsense in both directions — implausibly small burdens in the 1950s, negative burdens after 2008.

Longstaff (2011) provides the cleanest identification in the literature, and the paper's headline rests on it. He uses the municipal swap market: contracts exchanging a percentage of one-week taxable rates for tax-exempt index rates, observed weekly from August 2001 to June 2008. At one-week maturity, default and call risk are negligible, and the swap term structure lets the liquidity component be stripped. The naive short-end reading — municipal swap rates averaging 85.4 percent of taxable, implying a marginal tax rate of 14.6 percent in the working-paper version — becomes, after the correction, an identified marginal tax rate averaging 38.0 percent in the published version, which the paper uses throughout.⁴⁵

Two features make 38.0 percent credible as a market-clearing tax rate rather than an artifact. It exceeds the 35 percent top federal statutory rate that prevailed for most of the sample.

⁴ The published abstract states the marginal tax rate "averages 38.0%" — verified against the journal's own abstract metadata (Crossref, DOI 10.1111/j.1540-6261.2011.01650.x). The NBER working-paper version (w14687), verified line by line on file, reported 41.6 percent (its Table 3: mean 0.41574, s.d. 0.03784, range 0.31369-0.49826, N = 360); the naive 14.6 percent contrast and the sample dates are from the working-paper text. Table 4 carries the working-paper figure as a transparency row; pinpoint reconciliation against the published tables awaits library access.

⁵ Green (1993) supplies the theory of the relation between the taxable and tax-exempt curves when bonds need not be held to maturity; we cite it at claim level from secondary discussion, as the paper itself is paywalled at the time of writing.

That is consistent with a marginal holder pricing some state tax on the taxable alternative alongside the federal rate, and inconsistent with the folk claim that the marginal municipal buyer is always under-bracketed. The instrument imposes a caution, however: the identifying market is the variable-rate demand complex, whose dominant holders are tax-exempt municipal money funds holding national paper on which state exemption attaches only in-state, so the combined-rate interpretation is one reading of an above-statute rate, not a fact about households (verify instrument description). And in the working-paper version, Longstaff notes the average sits close to the top combined federal-plus-state rates of New York and California — 42.9 and 41.7 percent — the natural habitats of the marginal tax-exempt investor (verify pinpoints).

What the market charges is not what the federal government forgoes, and the distinction runs through every number that follows. The 38.0 percent is best read as a combined rate: the marginal investor prices the exemption against federal and, in part, state taxes. A correction to the federal schedule cannot coherently exceed the federal top rate. We therefore compute every headline object twice: at the full market rate, in which case the corrected schedule is explicitly the federal schedule augmented by the entire implicit tax, accruing across levels of government; and at a federal component capped at the 35 percent federal top rate of the sample period — a bound rather than an identification, since the split depends on deductibility and on which investors price the national index, choices the data do not settle (the 3.0-point residual is the state component under the simplest additive reading). At the capped rate the top-percentile correction is 0.30 points and the 2021 aggregate is \$44 billion.

Today's term structure is in Table 2, from three independent AAA curves published through the MSRB's public displays. At ten years the three providers land within little more than a point of one another: implied rates of 30.4, 30.3, and 29.3 percent, from curves as of August 6, 2026 against the most recent Treasury constant maturity (August 5, 2026). The implied rate falls to 22.7 percent at twenty years and 14.5 percent at thirty.⁶

⁶ The display curves are licensed for research use; their terms bar republication, so the published version of this paper will substitute a licensed Bloomberg BVAL pull by a coauthor, and the curve module of the replication package is written to be swapped. [WRDS robustness to be added: curves built directly from the MSRB transaction feed.]

Table 2. The term structure of the market-implied municipal tax rate, August 2026. AAA municipal curves from three providers via the MSRB's public displays (August 6, 2026); Treasury constant maturities from FRED (August 5, 2026). $\tau^* = 1 - (\text{municipal yield} / \text{Treasury yield})$, shown for the BVAL curve; the ICE column implies rates within roughly a point of BVAL's at each maturity, and the Tradeweb column differs by as much as three and a half points.

Maturity (years)	AAA muni, BVAL (%)	AAA muni, ICE (%)	AAA muni, Tradeweb (%)	Treasury CMT (%)	Implied τ^* , BVAL (%)
1	2.45	2.49	2.51	4.03	39.2
2	2.51	2.47	2.61	4.18	39.9
5	2.81	2.77	2.87	4.33	35.0
10	3.22	3.23	3.27	4.63	30.4
20	4.00	4.06	4.18	5.18	22.7
30	4.42	4.44	4.50	5.17	14.5

The Treasury benchmark is a choice, and it is the conservative one. The identification wants the default-free taxable twin of a AAA municipal curve; corporate AAA yields sit above Treasuries, so a corporate benchmark produces systematically higher implied rates — Poterba and Ramírez Verdugo (2011) report the corporate-based implicit rate at 29 percent in 2007 against 18.5 percent Treasury-based. Residual municipal credit and liquidity premia relative to Treasuries bias our measured τ^* down, understating the burden — and they are not small: Schwert (2017) attributes most of the tax-adjusted municipal spread to default risk (cited at claim level; the article is paywalled at the time of writing). Every known bias in the naive readings runs the same direction; the short-end identification exists precisely to strip them. The benchmarks also differ in which taxes they identify. Treasury interest is exempt from state income tax, so a ratio against Treasuries identifies a federal-only concept; the swap identification prices against state-taxable short rates and identifies the combined rate. Part of the term structure's apparent decline is benchmark, not maturity — and read as a federal-only concept, the one-year Treasury-based rate of 39.2 percent sits within two points of the 40.8 percent federal top rate on interest, a coherence check on the curve's own terms.

The term structure poses the question the rest of the paper must carry. The implied rate falls with maturity — from the 38.0 percent identified at Longstaff's one-week horizon to 14.5 percent at thirty years on today's curve. No investor's tax bracket depends on the maturity of the bond she holds. So either the long end is contaminated in ways even careful designs miss, or the long end is telling us something real about what the market prices.

It is not simply contamination. Chalmers (1998) examined municipal bonds secured by escrowed Treasury obligations — default-free, non-callable — and found that they still yield too

much relative to the tax story (cited from the discussions in Longstaff 2011 and Ang, Bhansali, and Xing 2010; the original is to be obtained before submission). The long-end puzzle survives clean credit controls, which rules out the easy dismissal.

Part of the answer is expectations. Longstaff's own estimates separate the current short rate from the long-run forward rate, and the forward marginal tax rate is 29.4 percent — the market prices the tax wedge as declining with horizon. Poterba (1986) made the point first: the taxable-tax-exempt spread embeds expected future tax policy, not just current rates. Kueng (2014) shows municipal yield spreads move with expected future tax rates at horizons out to thirty years. What this explains is why the long-end rate is lower, not that it is illusory: under the paper's own definition, burden is the yield the holder gives up, so a wedge narrowed by expectations of lower future rates is a smaller give-up for a purchaser today, not a larger one that the market has mismeasured. Expectations tell us the composition of the long-end rate — it is a forward-looking average of tax rates rather than a reading of today's — and that is exactly why the maturity grid in Table 4 is a grid of alternative burden rates and not a set of corrections to a single true one.

And part of the long end points the other way entirely. Ang, Bhansali, and Xing (2010) study municipal bonds trading at market discounts, whose price appreciation is taxed as ordinary income, and find implied tax rates of 70 to 100 percent — two to three times the statutory rate. Where the tax treatment of the marginal dollar is ordinary-income-like, the market capitalizes close to fully, at every maturity. The long end is not evidence that tax capitalization fades; it is evidence that what is being capitalized varies with the instrument.

The identification choice follows. The headline uses $\tau^* = 38.0$ percent because it is the cleanest credit- and liquidity-free identification available, a choice made on identification grounds; it sits near the top of the paper's robustness grid, and the text concedes as much wherever the headline appears. Its transplantation is a second caution: the estimate is from a 2001-2008 equilibrium, and no identified short-end estimate exists for 2021 — the intervening years brought the net investment income tax, the 2017 law's rate changes and state-tax deduction cap, and the collapse in bank and insurer municipal demand, any of which could move the marginal investor's rate. The 2021-contemporaneous variant is the 2021 ten-year rate of 35.6 in Table 4; today's short end (the one-year Treasury-based 39.2) suggests the neighborhood has not moved far. The alternative anchor — matching each holder's τ^* to a maturity — would require

maturity composition the household data do not contain and would import the unresolved puzzle into the headline. Instead the full term structure runs through every result: Table 4 reports the corrected schedule's key objects at 38.0, at the working-paper 41.6, at the federal component, at the 2021 ten-year, at Longstaff's long-run forward rate of 29.4, and at today's ten-, twenty-, and thirty-year rates.

The remaining objection is whether a one-week rate is informative about the burden on a household holding ten- and twenty-year bonds. The burden-relevant τ^* is the tax component of the yield give-up at whatever maturity the household holds. At the short end that component is isolated; at the long end it is entangled with the puzzle. Using 38.0 percent everywhere amounts to assuming the tax component is flat across maturities and assigning the rest of the long-end gap to non-tax factors; using the observed curve assumes the opposite. We do not claim the two bracket the truth. The long-end ratio contains several unidentified components whose signs are not uniformly established, and the short-end estimate is transported across maturity and across two decades, so neither end is demonstrably on the correct side of the true 2021 rate. The rows in Table 4 are therefore sensitivity scenarios — each a stated assumption carried through to its consequence — and what they establish is that the correction's sign and its concentration at the top of the income distribution do not depend on which assumption a reader prefers.

For the historical arc, no identified series exists before 2001. The long panel of Section 5's historical exhibit uses the Bond Buyer index because it is the only continuous municipal yield record back to the 1950s, labeled as the naive reading it is. From late 2016 the index's public series ends (the index itself continues behind the publisher's wall), and the figure splices on a AA quality curve — closer in credit quality to the index's 20-bond composite than a AAA curve — cross-checked against the AAA curve, which runs 20 basis points richer on average over the overlap panel. The 2020 excursion in the spliced segment is the real March 2020 dislocation, not a splice artifact.

4. Data and measurement

The household data are the 2022 Survey of Consumer Finances. The SCF is the only public United States microdata in which top-percentile municipal holdings are measurable: it oversamples wealthy households from an IRS-derived list frame, and its income questions cover calendar 2021, so every tax computation in the paper uses 2021 law. The survey's five implicates are pooled with weights the Board distributes already divided across implicates; the pooled weights sum to 131.3 million households.

Households are not tax units, and the conversion is not rebuilt here. We use the NBER SCF-to-TAXSIM files — produced by Kevin Moore's Federal Reserve Board conversion program and distributed by NBER — which split each SCF household into filing units with incomes, deductions, and dependents assigned by the program's rules. The analysis keeps the primary economic unit's tax units — 25,600 of them for 2022, representing 148 million weighted units — and excludes the financially independent co-residents whose assets the survey does not separate. That universe is smaller than the roughly 180 million units in return-based tables, which include dependent and other filers; levels here are therefore not directly comparable to published schedules, while differences within the paper are.

Statutory rates come from TAXSIM (Feenberg and Coutts 1993; version 35), which computes 2021 federal law including the net investment income tax and the alternative minimum tax. TAXSIM's default marginal rate is a wage-margin rate, and for municipal analysis it is wrong in both directions: a household's marginal rate on interest differs from its rate on wages whenever an extra dollar of interest crosses a threshold the wage margin does not. We therefore compute the interest margin directly, running every tax unit twice — once as reported, once with \$1,000 of additional taxable interest — and differencing the liabilities. A Massachusetts example from the state calculations illustrates the stakes: a couple whose default combined marginal rate reads 40.0 percent has a true interest margin of 43.7 percent, because the added interest also pushes capital gains across the net investment income tax threshold.

The reconstruction validates against its source. Our as-reported liabilities match the federal tax computed in the NBER files for 93 percent of units to within one dollar; the residual traces to TAXSIM revisions since those files were produced, concentrated in 2021's refundable credits, not to input mapping. Both of our runs use the same calculator on the same day, so the

differenced margins are internally consistent. Every submission to the calculator carries a pinned self-check case whose hand-computed liability must return exactly, and a run aborts if it does not (Appendix A1).

The municipal income measure is reported tax-exempt interest, which the SCF asks directly: \$81.4 billion in aggregate for 2021, held by 2.4 percent of tax units. The survey question spans interest received directly and through tax-exempt funds, matching the household-channel concept the aggregate anchors use. The alternative — imputing income as holdings times a current curve yield — is rejected because reported income captures the coupon flow on which the household actually bore the implicit tax, at the yields prevailing when its bonds were bought; imputation at 2021 curve yields, which were far below the stock's average coupon, would understate that flow by more than half — though reported coupons on premium bonds also exceed economic income, since part of each coupon amortizes the premium, a bias running the other way. The market-value construction — τ^* times the 2021 average ten-year Treasury yield of 1.44 percent times the \$2.36 trillion of survey holdings (direct bonds of \$792 billion plus tax-exempt funds of \$1,569 billion, the survey counterpart of the Financial Accounts household channel) — gives \$13 billion, the opportunity-cost burden at 2021 market terms. The two are not a bracket on one quantity. They measure different estimands — the give-up embedded in the coupon flow the household actually received, against the give-up on the stock it actually held at that year's market yields — and the survey cannot order them, because it records neither acquisition prices, nor vintages, nor premium amortization, nor maturity composition. Both appear in Section 5, each labeled with its own estimand, year and asset base. [WRDS robustness to be added: an imputation built on holdings-weighted vintages once maturity composition is available.]

The administrative record disagrees with the survey. IRS Statistics of Income tabulations report \$55.5 billion of tax-exempt interest for tax year 2021, on 6.6 million returns; the SCF aggregate exceeds it by a factor of 1.47 — and the gap has a second dimension: the survey has roughly half the holders (3.62 million units against 6.57 million returns) holding almost three times as much per holder (\$22,500 against \$8,400), so the survey's distributional shape may overstate concentration as well. Candidate explanations run in both directions: non-filing holders and under-reporting on line 2a deflate the SOI figure, while respondent rounding and gross reporting of fund distributions inflate the survey's. Rather than adjudicate, the paper scales: every

aggregate burden number has an SOI-scaled sibling, computed by multiplying by the ratio of the administrative to the survey base, and the distributional results — which depend on where tax-exempt interest sits, not on its level — are presented with the survey's shape in both cases.

Tax units are ranked by expanded income: cash income plus tax-exempt interest at the observed yield. The baseline effective rate is remitted federal individual income tax over expanded income; the corrected rate adds the implicit tax to both sides, per Section 2. The corrected object is the federal individual income tax schedule — not an all-federal-tax schedule of the kind CBO publishes; mapping the correction into a schedule that includes payroll and corporate taxes would dilute it by roughly the individual income tax's share of federal taxes at the top. Ranking is not driving the results: re-ranking units by corrected income moves the top-percentile correction from 0.34 to 0.35 points.

The public SCF has no state identifiers, and state taxes matter twice: they raise holders' statutory rates on taxable interest, and Section 3's identified market rate embeds them. The paper's headline is federal-only by design. As robustness, every municipal-holding unit is run through all fifty-one TAXSIM state calculators, and the results are combined two ways. A population-weighted expected state rate adds 4.5 percentage points to the dollar-weighted holder's marginal rate. Population weighting understates the add-on if municipal clientele concentrate in high-tax states holding in-state paper — a pattern documented directly by Babina, Jotikasthira, Lundblad, and Ramadorai (2021) — so an upper bound is also reported: assigning every holder her maximum-state rate adds 13.2 points, putting the dollar-weighted holder at 41.1 percent, above the identified market rate. Section 5 returns to what the pair establishes, and to why only the maximum-state figure is a bound in the literal sense.⁷

Aggregate anchors come from the Financial Accounts. As of the first quarter of 2026, \$4,456 billion of municipal securities are outstanding, of which 74.4 percent is held through channels that belong to taxable households: 46.5 percent directly, 18.5 percent through mutual funds, 6 percent through closed-end funds and exchange-traded funds, and 3.4 percent through money market funds. The SCF's direct-holdings aggregate of \$792 billion undercounts the Financial Accounts household line by roughly half on our 2022 comparison; Bergstresser and Cohen (2015) document the same undercount pattern in earlier waves. The paper therefore uses the Financial Accounts for stock-based aggregate constructions and the SCF for everything

⁷ TAXSIM's state law is frozen at 2021, which matches the income year exactly, so no overlay of later state changes is needed.

distributional.

Every number in the paper regenerates from public sources by one command. The replication package pulls the SCF, the NBER tax-unit files, the curves, the Financial Accounts and Treasury series, the SOI tables, and the Census state populations; runs the TAXSIM submissions with their pinned self-checks; and rebuilds every table, every figure, and a registry of every number quoted in the text, against which the prose is mechanically compiled. Appendix A5 gives the details.

The data cannot yet support three things. The public SCF's missing geography means state rates are expected rates, not actual ones; the restricted file would fix this. The household data contain no maturity composition — Section 5 states what that costs. And the curve history before 2001 is naive; [WRDS robustness to be added: author-constructed curves from the MSRB transaction feed, and MMD scale history, when institutional access arrives]. Each appears again in Section 8's forward design.

5. Results

Table 3 and Figure 1 report the paper's central object: average effective federal income tax rates by income percentile group, measured and corrected, for tax year 2021 at the anchor rate of 38.0 percent. The corrections are +0.34 percentage points for the 99th to 99.9th percentile (from 30.29 to 30.63 percent), +0.33 for the top 0.1 percent, +0.50 for the 95th to 99th, and +0.23 for the 90th to 95th. Below the 90th percentile no group moves by more than a tenth of a point. For all tax units together, the measured rate rises from 14.38 to 14.63 percent.⁸

⁸ Table entries are rounded from unrounded values and corrections are computed before rounding, so a displayed pair can differ from the displayed change by a hundredth. The marginal-rate column is a unit-weighted group mean; the text below reports dollar-weighted counterparts, which differ instructively.

Table 3. Effective federal income tax rates by income group, measured and corrected, tax year 2021. Tax units ranked by expanded income. The correction applies $\tau^* = 38.0$ percent to reported tax-exempt interest; Table 4 reports the variants. Source: authors' calculations from the SCF, TAXSIM, and the sources described in Section 4.

Income group	Tax units (m)	Share holding munis (%)	Tax-exempt interest (\$bn)	Avg MTR on interest (%)	ETR, measured (%)	ETR, corrected (%)	Change (pp)
0-20	29.6	0.1	0.1	4.8	-28.2	-28.2	+0.02
20-40	29.6	0.1	0.1	12.9	-7.5	-7.5	+0.01
40-60	29.6	1.1	0.4	16.0	0.1	0.1	+0.01
60-80	29.6	1.4	4.1	20.6	5.4	5.5	+0.09
80-90	14.8	4.1	3.9	27.5	11.8	11.9	+0.09
90-95	7.4	7.7	8.4	27.6	16.5	16.7	+0.23
95-99	5.9	19.4	32.9	35.6	22.4	22.9	+0.50
99-99.9	1.3	30.1	18.3	40.2	30.3	30.6	+0.34
Top 0.1	0.1	55.5	13.2	40.1	28.9	29.2	+0.33
All	148.1	2.4	81.4	16.8	14.4	14.6	+0.25




Figure 1. The correction is confined to the top of the distribution. Left: measured and corrected average effective federal income tax rates by income group. Right: the correction in percentage points. Tax year 2021, $\tau^* = 38.0$ percent.

Table 3's marginal-rate column reports the statutory side of the decomposition, and it rises exactly as the burden side requires: the average marginal rate on interest climbs from single digits in the bottom quintile to 27.5 percent in the 80th to 90th percentile group and just over 40 percent in the top percentile, where the 37 percent bracket and the net investment income tax bind

together. The measured column shows two things. The top 0.1 percent's measured effective rate sits below the 99th to 99.9th percentile's — the familiar composition effect of preferentially taxed capital gains at the very top — and the municipal correction, similar in size for both groups, does not close that gap. And the corrected column preserves the schedule's shape everywhere: the correction changes levels at the top, not orderings.

The geography of the correction follows the geography of the asset. Participation in tax-exempt interest runs from 2.4 percent of all tax units to 30.1 percent of the 99th to 99.9th percentile and 55.5 percent of the top 0.1 percent. Municipal income relative to total income peaks lower, in the 95th to 99th percentile group, which holds \$33 billion of the \$81.4 billion total; that is why the largest rate correction lands there rather than at the very top, where portfolios tilt toward equity and pass-through business income.

Tax year 2021 is a low baseline year, and the worry is quantitative rather than structural. Pandemic-era refundable credits pushed measured rates in the bottom half of the distribution to unusual lows — the bottom quintile's measured rate is deeply negative. The correction is not algebraically independent of them: it equals $I(Y - T)/[Y(Y + I)]$, which contains remitted tax T , and the credits move T . But it is numerically negligible wherever they bind, because I is negligible there: the largest movement below the 60th percentile is 0.02 points. The credits distort the baseline schedule the correction is applied to, not the correction.

Figure 2 reports the aggregate under each construction, stated with its base. Grossing up the survey's reported tax-exempt interest at the identified rate gives \$50 billion of unrecorded burden for tax year 2021; scaling to the administrative income base gives \$34 billion; the federal-component rate gives \$44 billion on the survey base. At current stocks and yields, applying the identified rate to the ten-year Treasury and the Financial Accounts household-channel stock gives \$58 billion a year, and the observed ten-year spread gives \$47 billion. The anchored construction deliberately crosses the short-end rate with the ten-year yield — an upper-bound pairing, labeled as such. At 2021 market terms the market-value construction gives \$13 billion — a different estimand, not the low end of a bracket: the give-up on the 2021 stock at 2021 yields, against the coupon-flow measure's give-up on the income actually received

(Section 4). For scale: calendar-2021 federal personal current taxes were \$2,128 billion, so the 2021 coupon-flow burden is 2.3 percent of that, and roughly twice Treasury's \$24 billion scored expenditure for public-purpose municipal interest — a comparison of adjacent but not identical bases, since reported tax-exempt interest includes private-activity bonds whose expenditures Treasury scores separately.⁹




Figure 2. The aggregate unrecorded burden under each construction. Sources and constructions as described in the text.

The rent decomposition yields the paper's most unexpected finding, and it survived two corrections that each moved it toward the same answer. The holder's counterfactual rate matters as much as the market's: a unit with a large municipal position and modest ordinary income shows a low marginal rate precisely because the munis shelter its taxable income, so the local rate on a small increment understates what the taxable twin would face. We therefore compute each holding unit's counterfactual rate two ways — the local rate on a \$1,000 increment, and the average rate on the full grossed-up block $y/(1 - \tau^*)$, from a third TAXSIM run. The block-average federal rate, dollar-weighted, is 34.3 percent against the local rate's 27.9; the gap is

⁹ The tax expenditure and the implicit tax measure different things: the expenditure scores forgone federal revenue given observed (depressed) yields, while the implicit tax measures the burden holders bear through those yields. They coincide only if the marginal and average holder's rates coincide.

the sheltering effect, concentrated exactly where the portfolios are largest (in the 95th to 99th percentile group the unit-weighted local mean is 35.6 percent but the dollar-weighted local rate is 22.2). Under the block-average rates and consistent tax-basis pairings, the aggregate 2021 rent at the anchor is approximately zero: plus \$1.0 billion on the expected-state pairing, minus \$0.9 billion on the federal-component pairing, plus \$12 billion only at the deliberately high maximum-state bound. Under the local rates the same three pairings run minus \$8.8, minus \$7.3, and plus \$4.1 billion, so the six consistent constructions together span minus \$8.8 to plus \$12.4 billion. Set against \$50 billion of burden, every one of them is small, and the two block-average pairings — the paper's preferred construction, because the block average is the rate a holder switching her whole position would face — sit within a billion dollars of zero. Measured at the short-end rate, the exemption delivered its average 2021 holder little or no aggregate windfall: the market priced the average municipal bond dollar closely.

The distribution behind the average carries more information than the average. 6.4 percent of municipal-holding tax units, weighted, have a measured federal marginal rate on interest of zero or below: on the federal margin alone, the exemption is worth nothing to them. Strict dominance is a stronger claim than the data support — it would need a taxable alternative matched on maturity, credit, liquidity and call terms, and the holder's full tax treatment including the state component the public survey does not carry, any of which can reverse the comparison. What the measure does establish is heterogeneity in measured federal interest margins, and a sensitivity in the aggregate. This is the modern echo of what Feenberg and Poterba (1991) found in the first tax-return data on tax-exempt interest. Candidate explanations are the familiar ones — advisor defaults, state-tax motives a federal margin misses, liquidity preferences, inertia after income falls — and the SCF cannot distinguish them. These units hold 16 percent of the survey's tax-exempt interest, and excluding them entirely lowers the 2021 aggregate from \$50 billion to \$42 billion — a real sensitivity, because the burden construction charges the market rate to every municipal dollar whatever the holder's own bracket. What the fact establishes is that the exemption's incidence within the holder population is far from uniform.

The anchor sits between the two state constructions. The dollar-weighted holder's combined local rate is 32.4 percent at the expected state rate and 41.1 at the maximum-state assignment; the block-average counterparts run about six points higher. The published anchor of 38.0 falls between them — the identification's marginal investor and the survey's largest holders

occupy the same neighborhood of combined rates, which is what near-zero aggregate rent requires. The claim is deliberately modest. Only the maximum-state figure is a bound in the literal sense, since no holder can face more than her own state's top rate; the population-weighted figure is an expectation under an assumption about where holders live, not a floor. What the pair establishes is consistency, not confirmation.

Rent flips sign across the curve, and that reframes the oldest question about the exemption. At the ten-year implied rate the top groups' rent is roughly zero; at the twenty- and thirty-year rates it is large and positive (the arithmetic follows from Table 4's rates and Section 2's formulas). The maturity structure of the implied rate is the rent structure of the exemption: short-maturity holders bear nearly their full statutory rate as implicit tax, while long-maturity holders keep most of theirs as windfall. The household data contain no maturity composition, so the paper reports the rates side by side rather than choosing among them — and the spread between them is itself the answer to whether the exemption is a windfall to the rich: it depends on which maturities they hold, a fact no distribution table has ever recorded either way. The rows are not a range containing the truth; each is what the rent would be if holders faced that rate. [WRDS robustness to be added: holdings-weighted maturity composition from the MSRB transaction feed.]

Figure 3's upper panel puts the naive implied rate against the top statutory rate on interest income from 1953 forward. The naive rate never approaches the statutory line: mid-century top rates of 70 to 91 percent coexisted with implied rates of 10 to 25 percent. That gap is the historical form of the windfall critique, and the figure is the case for taking the critique seriously — but the gap is not a measurement of historical rent. Rent is the distance between a holder's statutory rate and an identified market rate on a consistent basis, and the naive ratio is not one: it bundles credit, liquidity, call, maturity and benchmark effects with tax, which is the whole reason Section 3 rejects it for the modern period. No identified series exists before 2001, so the mid-century rent is unmeasured rather than large. The post-2008 sign flip described in Section 3 is visible at the right edge, and the spliced AA-curve segment from 2017 adds the March 2020 dislocation. The Longstaff correction, marked at the short end, is what an identified reading looks like against this backdrop.




Figure 3. Seven decades of the naive reading, and what identification changes. Upper panel: the naive implied rate (Bond Buyer 20-Bond over 20-year Treasury; AA curve splice from 2017, dashed) against the top statutory rate on interest income, with Longstaff's short-end identification marked. Lower panel: the implied burden series — the naive spread times the household-channel stock (the pathological series the paper rejects), and the anchored construction, $\tau^* = 38.0$ percent times the ten-year Treasury times the stock, from 2001.

The lower panel plots two burden series and does not compare them across eras. The naive-spread series repeats the upper panel's pathology in dollar form. The anchored series — the identified rate times the ten-year Treasury times the household-channel stock, the same construction as Figure 2 — runs from \$22 billion in 2001 to \$58 billion today. What the data do support historically is a statement about the base: household-channel municipal holdings grew from 4.7 percent of GDP in the 1950s to a peak of 20 percent in 2009, so whatever the true implicit tax rate was in the high-statutory-rate decades, the stock it applied to was several times smaller relative to the economy.

Table 4 collects the robustness grid. Across the τ^* variants — 38.0 published anchor, 41.6 working-paper version, 35.6 at the 2021 ten-year, 35.0 federal-component, 30.4 at today's ten-year, 29.4 at Longstaff's forward rate, 22.7 at twenty years, and 14.5 at thirty — the top-percentile correction runs from 0.09 to 0.39 points and the 2021 aggregate from \$14 billion to \$58 billion; at the published anchor the pair is 0.34 points and \$50 billion. The table's second pair of columns restates both quantities on the administrative base, which divides each by the survey-to-SOI ratio of 1.47: the correction at the published anchor becomes 0.23 points and the aggregate \$34 billion, and across the grid the correction runs 0.06 to 0.27 points. Appendix A3 gives the rest of the corrected schedule on that base and the near-linearity that makes the rescaling pass through. The state add-on moves the rent decomposition, not the federal correction. What survives every row is the correction's sign and its concentration above the 90th percentile; what the rows move is its size, a maturity call the data cannot yet settle.

Table 4. The correction under each τ^* variant, on both income bases. Correction to the 99th-to-99.9th percentile in percentage points and the 2021 aggregate, each reported on the survey base and on the administrative base — the survey rescaled so its tax-exempt interest aggregate matches IRS SOI. Ranks are the survey's on both bases (Appendix A3). Constructions as in Sections 2 and 4. Source: authors' calculations as in Table 3.

τ^* variant	τ^* (%)	P99–99.9 correction (pp), survey	P99–99.9 correction (pp), SOI base	Aggregate burden, 2021 (\$bn), survey	Aggregate burden, 2021 (\$bn), SOI base
Working-paper version of the same estimate	41.6	+0.39	+0.27	58	40
Short-end identified (Longstaff, published), 38.0%	38.0	+0.34	+0.23	50	34
10-year ratio, 2021 average curve	35.6	+0.31	+0.21	45	31
Federal component of the anchor (capped at the 35% federal top rate)	35.0	+0.30	+0.20	44	30
10-year ratio, August 2026	30.4	+0.24	+0.17	36	24
Longstaff's long-run forward rate	29.4	+0.23	+0.16	34	23
20-year ratio, August 2026	22.7	+0.16	+0.11	24	16
30-year ratio, August 2026	14.5	+0.09	+0.06	14	9

The paper stops short of three claims. It computes no Kakwani or Suits index: the corrected schedule is the primitive, and any index built on it inherits the correction mechanically. It makes no welfare statement and no reform counterfactual: those need incidence, which is Section 7's subject. And it attaches only range-style uncertainty to the corrections, because the

ranges show where the uncertainty lives: across the five SCF implicates the top-percentile correction runs 0.32 to 0.36 points, and moving τ^* one standard deviation of Longstaff's estimate in either direction moves it 0.29 to 0.40 — both narrow next to the 0.09-to-0.39 span of the τ^* grid itself. The rate choice dominates the sampling error; a full replicate-weight treatment belongs to the revision.

6. The consistency objection: housing and equity

Municipal bonds are not special, and this section takes up the objection that follows from that fact. Owner-occupied housing, retirement accounts, pass-through structures, life insurance — every preference that is capitalized into an asset's price drives the same wedge between remitted and borne. Correcting one asset held at the top of the distribution, and stopping there, biases the overall picture toward more measured progressivity. The objection is correct, and the paper's answer is not to dispute it.

The answer is a division by how much the market tells us, and the division is one of degree. Municipal bonds are the asset class where market prices come closest to identifying the implicit tax, because a tax-exempt bond and a taxable bond trade side by side and their yield ratio is observable every day. That observability is necessary and not sufficient: Section 3 showed the raw ratio is contaminated, the clean design is confined to a historical short-end market, and carrying its estimate to 2021 and across maturities is an assumption the Table 4 scenarios expose rather than resolve. Owner-occupied housing has no taxable twin trading beside it at all; corporate equity has a marginal investor who may not be taxable at all. So the paper corrects the asset the market speaks to, sizes the two large neighbors with explicit arithmetic, and states the direction each pushes the headline. The arithmetic below is quantitative rather than rhetorical, but it is a scale illustration under stated assumptions — not a bound, and not a correction: neither exercise brackets a true value, and this section does not produce a percentile-level schedule for either asset.

6.1 Housing

The housing correction has a known direction, which is less than a known size. The preference bundle — untaxed imputed rent, the mortgage interest deduction, deductible property taxes — is capitalized into house prices: Poterba (1984) is the asset-market treatment, the user-cost framework of Himmelberg, Mayer, and Sinai (2005) carries it, and Sommer and Sullivan (2018) find in general equilibrium that repealing the mortgage interest deduction lowers house prices by about 4.2 percent. Capitalization means owner-occupants earn a depressed implicit return on housing as bondholders earn a depressed yield on munis. The analogy stops short of an equivalence, and the gap matters for what can be computed: capitalization into the purchase price supports a lower housing-service yield at current market prices, but it does not by itself deliver an annual household-level rate correction, and it leaves the transitional gains and losses — which accrued to whoever owned when the preference was enacted or changed — unallocated across cohorts. Housing is held far further down the distribution: 80.9 percent of households between the 50th and 90th percentiles own homes worth \$18.2 trillion, against \$0.21 trillion of municipal holdings in the same group.

The magnitudes are what make the question live. Treasury scores the three housing preferences at \$189 billion a year for fiscal 2024 — the imputed-rent exclusion alone is among the largest individual income tax expenditures — against \$24 billion for public-purpose municipal interest: roughly eight to one. If a quarter of the housing preference were capitalized into depressed implicit returns, the implied middle-class implicit tax would be of the same order as the entire municipal correction, and it would land on the 50th to 90th percentiles rather than the top. State what that calculation is and is not. The one-quarter conversion is assumed, not derived; an annual tax-expenditure score is a revenue estimate rather than a measure of depressed annual returns, and the step from one to the other is exactly the capitalization model this paper does not build. So the arithmetic establishes that the housing margin is large enough to matter and that it runs the other way — not that it would offset most of the municipal correction, which is a claim this paper cannot support and does not make. One nuance disciplines it further: imputed rent is absent from the income base of the CBO, Treasury, and JCT frameworks altogether, so for

them both the income and its implicit tax are missing and the bias runs through the denominator rather than the recorded rate; the symmetry argument as stated applies to the national-accounts-based frameworks that do include imputed rent. The direction — a correction landing in the middle of the distribution — survives under either treatment. A second nuance cuts the other way: the municipal wedge is collected by a government as borrowing-cost revenue, which is Section 2's reason an avoidable wedge counts as a tax; the housing wedge is capitalized into prices and accrued to whoever owned at enactment, with no government collecting it. The two corrections are symmetric as measurement of depressed pre-tax returns, not as taxation — which weakens the consistency objection's force even as the paper concedes its direction.¹⁰

6.2 Equity

A benchmark arithmetic poses the equity question. Suppose dividend and capital-gains preferences depress pre-tax equity returns by twenty basis points. On \$91.9 trillion of domestic corporate equity that is \$184 billion a year, several times the municipal numbers. The arithmetic is right; the two premises it needs are not established. The base matters by a factor of nearly four: only 27 to 28 percent of United States corporate equity sits in taxable accounts, with roughly 32 percent of publicly traded equity held by foreigners and 34 percent in retirement accounts (Rosenthal and Mucciolo 2024). And the existence of any aggregate wedge at all depends on who the marginal investor is.

That dependence is the marginal-investor problem, and it is fatal to a point estimate rather than merely attenuating. If the price-setting investor in the aggregate equity market is a pension fund or a foreign institution, individual-level dividend and capital-gains preferences are not capitalized into prices at all: the implicit tax on equity is zero, and the preferences are pure windfall to taxable holders. Whether that is the case is exactly what the empirical literature splits on.

The split is real and the paper preserves it. Sialm (2009) estimates an aggregate effective tax yield of about 0.4 percent of equity value in 2006 — roughly forty basis points, and itself the low end of a series that has declined from over 3.5 percent in 1950 as rates fell and ownership

¹⁰ Poterba and Sinai (2008) give the distributional texture within owner-occupancy: user costs of 5.0 percent for the highest-income owners against 6.9 percent below \$40,000, and mortgage-interest tax savings that rise steeply with income. The explicit expenditures favor the top; the capitalized implicit tax reaches the middle.

migrated to untaxed accounts. Amromin, Harrison, and Sharpe (2008) find no aggregate imprint of the 2003 dividend tax cut on stock prices. Auerbach and Hassett (2005) find cross-sectional price responses to the same reform; Chetty and Saez (2005) find payout responses. The reference points stand at zero and roughly forty basis points, and we refuse to average them. The scale check a referee will run: forty basis points on the \$24.8 trillion of taxable-account equity is \$99 billion a year — larger than the entire municipal correction. Note what the upper figure is. An effective tax yield measures tax paid relative to equity value; it does not identify how much of that tax is capitalized into equilibrium pre-tax returns, and the relative-price evidence below does not identify that level either. Forty basis points is therefore a conditional scenario, not a demonstrated ceiling — the true aggregate wedge could sit above it as easily as between it and zero. An equity margin of this scale, with its sign unsettled and its magnitude unbounded in either direction, can swallow the headline, which is precisely why the paper's quantitative claims are munis-only. Two questions sharpen the range, and the next paragraphs answer them in turn: does modern ex-dividend-day evidence show dividend taxes reaching prices at all, and do partnership structures show the capital-gains preference being priced?

The ex-dividend-day literature is the sharpest lens on whether dividend taxes reach prices, and its modern arc answers the first question. Elton and Gruber (1970) made ex-day pricing a tax thermometer: on 1966-67 data the ex-day price drop implied a marginal stockholder bracket of 36.4 percent, matching tax-return-based estimates, with clientele patterns by yield. The modern reforms supplied the experiments, and they produced both results at once. Zhang, Farrell, and Brown (2008) find the mean price-drop ratio rose from 0.65 before the cut to 0.86 after, significant at the 1 percent level, with tax-motivated ex-day volume receding. Bell and Jenkinson (2002) find the United Kingdom's 1997 removal of pension funds' dividend credits moved the valuation of high-yield companies — identifying pension funds, not top-bracket individuals, as the marginal investor in that market.

The same literature explains why those results cannot size an aggregate wedge. Chetty, Rosenberg, and Saez (2005) show the 2003 comparison fails placebo tests: the relation between tax rates and ex-day premia over 1962-2004 is so noisy that responses cannot be credibly detected even around large tax reforms. Hartzmark and Solomon (2019) document non-tax demand for dividends that contaminates ex-day pricing. The answer to the first question, then: modern ex-dividend-day studies show dividend taxes in relative prices where taxable traders are

pivotal, and they neither reject the zero reference point nor the forty-basis-point one for the aggregate. The two stay where they are, unranked.

Partnership structures carry the capital-gains preference, and they answer the second question in the affirmative, with a caveat that leaves the aggregate as unsettled as before. Individual capital-gains taxes demonstrably move equity prices: in the week of the May 1997 budget accord that cut the top rate, non-dividend-paying stocks returned about 12 percent against 6 percent for payers (Lang and Shackelford 2000), and Dai, Maydew, Shackelford, and Zhang (2008) identify both capitalization and lock-in effects around the same reform, sorted by taxable ownership. Where a partnership structure packages a tax advantage for an identifiable taxable clientele, the advantage is priced at close to its arithmetic value. Sinai and Gyourko (2004) find the 1997 rate cut shrank the UPREIT partnership deferral subsidy: acquisitive UPREITs fell roughly 14 percent relative to non-acquisitive UPREITs, in a design benchmarked against regular REITs. In 2021, Magellan Midstream's board publicly priced the present value of corporate taxes from converting its master limited partnership at about \$2.3 billion and found no offsetting valuation premium from corporate form: the pass-through shield is in the price. Congress restricted publicly traded partnerships in 1987 precisely because the form eroded the corporate tax base (Congressional Research Service 2011).

The caveat: every clean partnership result is a relative price — UPREIT against REIT, MLP against C-corporation, high-yield against low-yield stock. Relative-price capitalization confirms the mechanism this paper applies to municipal bonds: where the taxable investor is pivotal, her taxes are in the price. It does not identify the level of aggregate pre-tax C-corporation returns against a no-tax counterfactual, which is what an economy-wide equity correction needs. The margin is quantitatively first-order — pass-throughs now generate over half of United States business income, far more concentrated at the top than C-corporation income (Cooper et al. 2016) — so the question matters; it is simply not answerable from relative prices. The answer to the second question: partnership evidence moves the plausibility of equity capitalization up, and the aggregate wedge remains unidentified between the corners.

What a full treatment would require, stated once: for housing, a capitalization model joining imputed rent, deductions, and prices; for equity, a marginal-investor identification that does not currently exist; for retirement and insurance vehicles, deferral accounting. Section 8 assigns them to the research design. This paper does one asset from market prices, and for the

other two states the direction, the scale, and the assumption each figure rests on.

7. The incidence question

Everything to this point is partial equilibrium, and the assumption should be stated as an assumption. The correction takes observed prices — observed municipal yields, observed taxable yields, the observed gap — and reallocates measured burden at those prices. It makes no claim about what yields would be in a different tax system, and it needs none: the implicit tax is a burden current law imposes and current prices reveal. What the general-equilibrium question governs is not the measurement but its interpretation — who ultimately bears the burden, and who collects it.

Who collects is the easy half. The implicit tax is real revenue: state and local issuers collect it as lower borrowing costs on roughly \$4,456 billion of debt. The federal distribution table records neither side of the transaction — not the levy on the holder, not the receipt by the issuer. Galper, Rosenberg, Rueben, and Toder (2013) distributed the receipt side: under the traditional allocation the top 1 percent receive 57 percent of the exemption's benefit, and their capital-market simulation shifts part of it to issuers; this paper's correction is the levy side of the same transaction. Together they describe a tax system operating entirely off the federal books.

The rent term connects the correction to the oldest efficiency critique of the exemption. Gordon and Slemrod (1983) modeled the exemption as subsidizing municipal borrowing through interest-rate adjustment; the inframarginal holder's rent is the leakage in that mechanism — the part of the federal revenue cost that reaches neither the issuer nor any public purpose. The decomposition prices that leakage at the household level: approximately zero at the short end, most of the statutory rate at the long end (Section 5). The efficiency critique, like the windfall critique, turns out to be a statement about maturity.

The equilibrium behind the intuition has a name. Miller (1977) described the capital-market equilibrium in which pre-tax interest rates gross up until the marginal taxable investor is indifferent between taxable and tax-exempt returns (cited at claim level; the article is to be obtained). In that world, every fully taxed lender is compensated for her taxes through the grossed-up rate; the tax on interest is borne by borrowers; and tax-exempt lenders — endowments, pension funds, foreign official institutions — collect the gross-up as pure rent. The question at this project's origin — whether pre-tax interest rates are higher because of taxes, with

the difference accruing to untaxed institutions while borrowers bear it — is the Miller equilibrium stated in one sentence.

The strongest empirical statement in that direction is aggregate. Gordon and Slemrod (1988), on 1983 data, estimated that the United States collected approximately zero net revenue from taxing capital income once arbitrage across brackets and instruments is netted out: high-bracket taxpayers deduct interest at high rates while low-bracket and exempt lenders report it at low ones. Their counterfactuals raised revenue "mainly at the expense of those in the highest tax brackets." Revisiting the question on 1995 data, Gordon, Kalambokidis, and Slemrod (2004) found the net tax on capital income had turned positive — roughly \$108 billion — so the near-zero result describes the 1980s tax system, not a constant of nature. To whatever extent capital income taxation shuffles rather than collects, distribution tables built on remittance misallocate who pays for it — the aggregate form of this paper's percentile-level point — and the shuffle was largest exactly when statutory rates were highest.

The open-economy literature points the opposite way, and the two are not reconcilable by assumption. What world capital markets pin is the after-tax return investors require, not the domestic pre-tax rate (Gordon 1986; Razin and Sadka 1991). A source-based capital income tax therefore drives capital out until the domestic pre-tax marginal return has risen enough to restore that required return: the pre-tax rate moves, and the burden comes to rest on the immobile factors whose returns fall — which is why Gordon's own account has countries that tax capital more heavily paying higher real interest rates, and why Razin and Sadka find the optimal tax on perfectly mobile capital vanishing. The consequence for this paper is narrower than a fixed rate level would give. There is still no economy-wide transfer to tax-exempt lenders of the Miller kind, because the pre-tax rate here rises to compensate mobile capital rather than to compensate taxed domestic lenders; but the exemption's general-equilibrium footprint is not confined to relative prices either, and this stance does not by itself rule a gross-up out. Fullerton and Rogers (1991; 1993) remain the most serious attempt to trace tax burdens through factor prices in general equilibrium, over lifetimes rather than years, and their machinery would be needed to adjudicate; nothing lighter will.

The direct evidence that would choose between the stories barely exists. The closest sustained attempt is the tax-adjusted Fisher-effect literature — Darby (1975) and Feldstein (1976) derived the gross-up that nominal rates should exhibit, and the empirical work that

followed largely failed to find it (Summers 1983) — results whose negativity supports this section's agnosticism. Event studies of tax changes on equity prices exist in quantity; clean estimates of tax capitalization into the level of pre-tax interest rates do not. We searched for a credible design for identifying the interest-rate gross-up — both to cite and to propose — and found none; the central empirical question of this section is open.

One implication of the Miller story has never been quantified by anyone, and this paper does not quantify it either. If pre-tax rates embed any gross-up, the tax-exempt sector collects it on its entire taxable fixed-income portfolio — every endowment, every pension fund, every foreign official holder of Treasuries and corporates. Measuring the transfer requires the sector's taxable holdings by class, an equilibrium gross-up per incidence stance, and a defended stance. Each piece is named in Section 8's design; none is delivered here.

What each stance does to this paper's numbers can be stated exactly. Under open-economy pinning of the required after-tax return, the correction stands as measured — the yields in it are observed yields — and the ultimate burden lands on immobile domestic factors rather than on the holder. Under a full Miller gross-up, the correction stands as measured — prices are still observed prices — and a second, institution-side table is missing: the rent collected by tax-exempt lenders, which no framework records either. Under neither stance is the remittance convention's zero the right entry for the municipal holder. The incidence stances disagree about where the burden ultimately comes to rest, not about whether the holder's measured rate is wrong.

The wedge between remitted and borne at observed prices is invariant to the incidence stance. What it ultimately means for the distribution of burdens awaits evidence that does not yet exist, and the paper leaves the question open rather than assuming it closed.

8. Conclusion

This paper measures what every maintained distribution table leaves out for the one asset where market prices permit the measurement. Municipal bondholders bear an implicit tax through depressed pre-tax yields; the tables record their burden as zero. Combining SCF portfolios, TAXSIM statutory rates, and a market-identified implicit tax rate, the paper produces corrected effective federal income tax rate schedules by percentile and reports the burden and the inframarginal rent separately.

Three statements survive every construction. Top-percentile effective rates are understated — by 0.34 percentage points at the identified market rate, 0.30 at its federal component, never less than 0.09 anywhere on the paper's grid of market-implied rates — with an aggregate unrecorded burden between \$47 and \$58 billion a year at current stocks and yields. The size of the correction is a maturity question, because the term structure of the market-implied rate is genuinely open. And at the short end the exemption delivered essentially no aggregate windfall in 2021: under the two pairings whose counterfactual rates are computed on the full taxable block, the rent is within a billion dollars of zero either side, and across every consistent pairing it stays small against the burden. What that establishes is where the rent is not. Estimated rent is near zero at the short-end anchor and rises as the assumed rate falls toward the levels observed further out the curve — which is a statement about the paper's scenarios, not a measurement of what long-maturity holders were actually paid in 2021, a quantity the household data cannot locate.

The correction carries its own bias. Correcting municipal bonds alone tilts measured progressivity upward. The housing preferences are eight times larger in scored expenditure and sit in the middle of the distribution; capitalized even partially, they push the other way. The equity margin has reference points at zero and at an amount larger than the entire municipal correction, with no identification that would place the truth between them. The headline is not usable without those two neighbors stated alongside it.

The forward design is concrete. Author-constructed curves from the MSRB transaction feed and the MMD scale history will replace the display curves and extend the identified record [WRDS robustness to be added]. The restricted SCF's geography will turn expected state rates into actual ones. Holdings-weighted maturity composition will replace the term-structure scenarios with a rate matched to what households actually hold. The institution-side table — the transfer to tax-exempt lenders under each incidence stance — is the natural companion paper. And the full program is the one Galper and Toder began four decades ago: every capitalized preference, both sides of the ledger, with market identification where markets exist.

The distribution tables are not wrong to count remittances; they are incomplete in a way that market prices can now partly repair. This paper makes the repair where the market permits it, states the direction and the scale of what it cannot repair, and reports both the burden the tables hide and the windfall they hide — so that the next argument about who pays can start from a

schedule that includes what the municipal bondholder actually gives up.

Appendix

A1. Constructions and validation

The after-tax-income invariance of Section 2: baseline after-tax income is $Y - T$. The corrected schedule replaces Y with $Y + y \cdot \tau^* / (1 - \tau^*)$ and T with $T + y \cdot \tau^* / (1 - \tau^*)$; the difference of the two is unchanged. The correction moves only the ratio.

The TAXSIM submissions run in chunks of one thousand tax units. Every chunk is prepended with a pinned self-check unit and the pipeline aborts if the calculator returns anything else for it. The unit is a 2021 single filer with \$50,000 of wages and \$100 of taxable interest, whose federal liability of \$2,907.00 is verified by hand: the \$12,550 standard deduction leaves \$37,550 of taxable income, the 10 and 12 percent brackets give \$4,307.00, and the \$1,400 recovery rebate credit — full, at an income well below the phase-out — nets \$2,907.00. The pin is deliberately on the analysis year and deliberately in the credit region: a check on a neighboring year would exercise the submission and response handling without exercising target-year parameter selection, and a check on wage tax alone would not touch the refundable credits that shape this paper's 2021 baseline schedule. The two-run interest margin uses an increment of \$1,000 of taxable interest. Against the federal liabilities shipped with the NBER tax-unit files, our as-reported run matches 93 percent of units within one dollar; inspection of the residual cases shows differences concentrated in 2021 refundable credits, consistent with calculator revisions after those files were produced.

A2. The expected-state-rate calculation

Every municipal-holding tax unit is run through all fifty-one TAXSIM state calculators (the fifty states and the District of Columbia) under 2021 law, with the same two-run interest differencing. The population-weighted expectation uses 2021 Census state population estimates. The dollar-weighted expected state marginal rate on interest across holders is 4.5 percent; the maximum-state bound is 13.2 percent. TAXSIM's state law is fixed at 2021, which matches the income year.

A3. The robustness grid, extended

Table 4 reports the grid on both bases. Rescaling the survey's tax-exempt interest to the administrative record multiplies every aggregate by 0.68: the 2021 aggregate at the identified rate becomes \$34 billion and the federal-component aggregate \$30 billion.

The percentile corrections scale too. A group's correction is $B(Y - T)/[Y(Y + B)]$ with $B = y \cdot \tau^*/(1 - \tau^*)$, which at these magnitudes is very nearly linear in y , so a proportional rescaling of tax-exempt interest passes through almost exactly: the 99th-to-99.9th percentile's +0.34 becomes +0.23, the 95th-to-99th's +0.50 becomes +0.34, and the top 0.1 percent's +0.33 becomes +0.23. What proportional scaling leaves unchanged is the shape — the ordering of groups by correction size, and the confinement of any material correction above the 90th percentile — not the levels.

Ranks are held at the survey's on both bases. A uniform factor is a statement about the aggregate and carries no unit-level information about which respondents over-report, so re-ranking units on rescaled income would reshuffle them on the strength of an assumption rather than of data, and would move the thinnest cells most.

A4. The survey-administrative gap

The SCF reports \$81.4 billion of tax-exempt interest for 2021; SOI reports \$55.5 billion on 6.6 million returns. The ratio is 1.47. The count dimension sharpens it: 3.62 million survey holder-units against 6.57 million administrative returns, with survey holders reporting \$22,500 each against \$8,400 — the survey misses small holders, inflates large ones, or both, which bears on the distributional shape and not just the level. A by-AGI-class shape comparison against SOI Table 1.4 (already in the replication pipeline) is the revision's first empirical task. The gap's candidate components, none of which the public data can size: non-filing holders (deflates SOI), under-reporting of line 2a information items (deflates SOI), respondent rounding and recall (inflates SCF), and gross reporting of fund distributions (inflates SCF). Both bases run through every aggregate in the paper.

A5. Reproducibility

One command regenerates the paper's every number, table, and figure from public sources: the Federal Reserve Board's SCF files, the NBER SCF-to-TAXSIM tax-unit files, the TAXSIM35 web calculator, the municipal AAA curves from the MSRB's public displays, FRED (Treasury yields, the Bond Buyer historical index, the Financial Accounts, NIPA receipts), IRS SOI Table 1.4, Census state population estimates, and Treasury's tax-expenditure tables. Intermediate artifacts are cached; the TAXSIM runs carry the pinned self-check described in A1; and the text of the paper is compiled against a registry of pipeline outputs, so a quoted number cannot drift from the scripts that produce it. The replication package accompanies the draft.

References

- Amromin, Gene, Paul Harrison, and Steven Sharpe. 2008. "How Did the 2003 Dividend Tax Cut Affect Stock Prices?" *Financial Management* 37(4): 625–646.
- Ang, Andrew, Vineer Bhansali, and Yuhang Xing. 2010. "Taxes on Tax-Exempt Bonds." *Journal of Finance* 65(2): 565–601.
- Auerbach, Alan J., and Kevin A. Hassett. 2005. "The 2003 Dividend Tax Cuts and the Value of the Firm: An Event Study." NBER Working Paper 11449.
- Auten, Gerald, and David Splinter. 2024. "Income Inequality in the United States: Using Tax Data to Measure Long-Term Trends." *Journal of Political Economy* 132(7): 2179–2227.
- Babina, Tania, Chotibhak Jotikasthira, Christian Lundblad, and Tarun Ramadorai. 2021. "Heterogeneous Taxes and Limited Risk Sharing: Evidence from Municipal Bonds." *Review of Financial Studies* 34(1): 509–568.
- Bell, Leonie, and Tim Jenkinson. 2002. "New Evidence of the Impact of Dividend Taxation and on the Identity of the Marginal Investor." *Journal of Finance* 57(3): 1321–1346.
- Bergstresser, Daniel, and Randolph Cohen. 2015. "Changing Patterns in Household Ownership of Municipal Debt: Evidence from the 1989–2013 Surveys of Consumer Finances." Brandeis University Economics Working Paper 87.
- Chalmers, John M. R. 1998. "Default Risk Cannot Explain the Muni Puzzle: Evidence from Municipal Bonds That Are Secured by U.S. Treasury Obligations." *Review of Financial Studies* 11(2): 281–308.
- Chetty, Raj, Joseph Rosenberg, and Emmanuel Saez. 2005. "The Effects of Taxes on Market Responses to Dividend Announcements and Payments: What Can We Learn from the 2003 Dividend Tax Cut?" NBER Working Paper 11452.
- Chetty, Raj, and Emmanuel Saez. 2005. "Dividend Taxes and Corporate Behavior: Evidence from the 2003 Dividend Tax Cut." *Quarterly Journal of Economics* 120(3): 791–833.
- Congressional Budget Office. 2024. *The Distribution of Household Income in 2021*. Washington: CBO.
- Congressional Research Service. 2011. "Master Limited Partnerships: A Policy Option for the Renewable Energy Industry." Report R41893.
- Cooper, Michael, John McClelland, James Pearce, Richard Prisinzano, Joseph Sullivan, Danny Yagan, Owen Zidar, and Eric Zwick. 2016. "Business in the United States: Who Owns It, and How Much Tax Do They Pay?" *Tax Policy and the Economy* 30: 91–128.
- Cronin, Julie-Anne. 2022. "U.S. Treasury Distributional Analysis Methodology." Office of Tax Analysis Technical Paper 8 (update).
- Dai, Zhonglan, Edward Maydew, Douglas A. Shackelford, and Harold H. Zhang. 2008. "Capital Gains Taxes and Asset Prices: Capitalization or Lock-in?" *Journal of Finance* 63(2):

709–742.

- Darby, Michael R. 1975. "The Financial and Tax Effects of Monetary Policy on Interest Rates." *Economic Inquiry* 13(2): 266–276. (verify)
- Elton, Edwin J., and Martin J. Gruber. 1970. "Marginal Stockholder Tax Rates and the Clientele Effect." *Review of Economics and Statistics* 52(1): 68–74.
- Feenberg, Daniel R., and Elisabeth Coutts. 1993. "An Introduction to the TAXSIM Model." *Journal of Policy Analysis and Management* 12(1): 189–194.
- Feenberg, Daniel R., and James M. Poterba. 1991. "Which Households Own Municipal Bonds? Evidence from Tax Returns." *National Tax Journal* 44(4): 93–103.
- Feldstein, Martin. 1976. "Inflation, Income Taxes, and the Rate of Interest: A Theoretical Analysis." *American Economic Review* 66(5): 809–820. (verify)
- Fullerton, Don. 1984. "Which Effective Tax Rate?" *National Tax Journal* 37(1): 23–41.
- Fullerton, Don, and Diane Lim Rogers. 1991. "Lifetime Versus Annual Perspectives on Tax Incidence." *National Tax Journal* 44(3): 277–287.
- Fullerton, Don, and Diane Lim Rogers. 1993. *Who Bears the Lifetime Tax Burden?* Washington: Brookings Institution.
- Galper, Harvey, and Eric Toder. 1984. "Transfer Elements in the Taxation of Income from Capital." In Marilyn Moon, ed., *Economic Transfers in the United States*, 87–138. Chicago: University of Chicago Press for NBER.
- Galper, Harvey, Joseph Rosenberg, Kim Rueben, and Eric Toder. 2013. "Who Benefits from Tax-Exempt Bonds? An Application of the Theory of Tax Incidence." Tax Policy Center working paper.
- Gordon, Roger H. 1986. "Taxation of Investment and Savings in a World Economy." *American Economic Review* 76(5): 1086–1102.
- Gordon, Roger H., Laura Kalambokidis, and Joel Slemrod. 2004. "Do We Now Collect Any Revenue from Taxing Capital Income?" *Journal of Public Economics* 88(5): 981–1009.
- Gordon, Roger H., and Joel Slemrod. 1983. "A General Equilibrium Simulation Study of Subsidies to Municipal Expenditures." *Journal of Finance* 38(2): 585–594.
- Gordon, Roger H., and Joel Slemrod. 1988. "Do We Collect Any Revenue from Taxing Capital Income?" *Tax Policy and the Economy* 2: 89–130.
- Green, Richard C. 1993. "A Simple Model of the Taxable and Tax-Exempt Yield Curves." *Review of Financial Studies* 6(2): 233–264.
- Hartzmark, Samuel M., and David H. Solomon. 2019. "The Dividend Disconnect." *Journal of Finance* 74(5): 2153–2199.

- Himmelberg, Charles, Christopher Mayer, and Todd Sinai. 2005. "Assessing High House Prices: Bubbles, Fundamentals and Misperceptions." *Journal of Economic Perspectives* 19(4): 67–92.
- Iglarsh, Harvey J., and Ronald Gage Allan. 2003. "Implicit Taxes and Progressivity." In Thomas M. Porcano, ed., *Advances in Taxation* 15. Bingley: Emerald. (engaged from the published abstract only)
- Joint Committee on Taxation. 2013. "Modeling the Distribution of Taxes on Business Income." JCX-14-13.
- Kueng, Lorenz. 2014. "Tax News: Identifying Tax Expectations from Municipal Bonds with an Application to Household Consumption." NBER Working Paper 20437.
- Lang, Mark H., and Douglas A. Shackelford. 2000. "Capitalization of Capital Gains Taxes: Evidence from Stock Price Reactions to the 1997 Rate Reduction." *Journal of Public Economics* 76(1): 69–85.
- Longstaff, Francis A. 2011. "Municipal Debt and Marginal Tax Rates: Is There a Tax Premium in Asset Prices?" *Journal of Finance* 66(3): 721–751.
- Miller, Merton H. 1977. "Debt and Taxes." *Journal of Finance* 32(2): 261–275.
- Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2018. "Distributional National Accounts: Methods and Estimates for the United States." *Quarterly Journal of Economics* 133(2): 553–609.
- Poterba, James M. 1984. "Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach." *Quarterly Journal of Economics* 99(4): 729–752.
- Poterba, James M. 1986. "Explaining the Yield Spread between Taxable and Tax-Exempt Bonds: The Role of Expected Tax Policy." In Harvey S. Rosen, ed., *Studies in State and Local Public Finance*, 5–52. Chicago: University of Chicago Press for NBER.
- Poterba, James M., and Arturo Ramírez Verdugo. 2011. "Portfolio Substitution and the Revenue Cost of the Federal Income Tax Exemption for State and Local Government Bonds." *National Tax Journal* 64(2): 591–614.
- Poterba, James M., and Todd Sinai. 2008. "Tax Expenditures for Owner-Occupied Housing: Deductions for Property Taxes and Mortgage Interest and the Exclusion of Imputed Rental Income." *American Economic Review* 98(2): 84–89.
- Razin, Assaf, and Efraim Sadka. 1991. "Vanishing Tax on Capital Income in the Open Economy." NBER Working Paper 3796.
- Rosenthal, Steven M., and Livia Mucciolo. 2024. "Who's Left to Tax? US Taxation of Corporations and Their Shareholders." Tax Policy Center.
- Saez, Emmanuel, and Gabriel Zucman. 2024. "Distributional Tax Analysis in Theory and Practice: Harberger Meets Diamond-Mirrlees." Working paper.

- Scholes, Myron S., and Mark A. Wolfson. 1992. *Taxes and Business Strategy: A Planning Approach*. Englewood Cliffs: Prentice-Hall.
- Schwert, Michael. 2017. "Municipal Bond Liquidity and Default Risk." *Journal of Finance* 72(4): 1683–1722.
- Sialm, Clemens. 2009. "Tax Changes and Asset Pricing." *American Economic Review* 99(4): 1356–1383.
- Sinai, Todd, and Joseph Gyourko. 2004. "The Asset Price Incidence of Capital Gains Taxes: Evidence from the Taxpayer Relief Act of 1997 and Publicly-Traded Real Estate Firms." *Journal of Public Economics* 88(7): 1543–1565.
- Slemrod, Joel. 2008. "Does It Matter Who Writes the Check to the Government? The Economics of Tax Remittance." *National Tax Journal* 61(2): 251–275. (verify)
- Sommer, Kamila, and Paul Sullivan. 2018. "Implications of US Tax Policy for House Prices, Rents, and Homeownership." *American Economic Review* 108(2): 241–274.
- Splinter, David. 2020. "U.S. Tax Progressivity and Redistribution." *National Tax Journal* 73(4): 1005–1024.
- Summers, Lawrence H. 1983. "The Non-Adjustment of Nominal Interest Rates: A Study of the Fisher Effect." In James Tobin, ed., *Macroeconomics, Prices and Quantities: Essays in Memory of Arthur M. Okun*. Washington: Brookings Institution. (verify)
- Williams, Michael. 2002. "Tax Clienteles and Implicit Taxes: Interaction of Multiple Clienteles in Financial Markets." Working paper, Anderson Graduate School of Management, UCLA.
- Zhang, Yi, Kathleen A. Farrell, and Todd A. Brown. 2008. "Ex-Dividend Day Price and Volume: The Case of the 2003 Dividend Tax Cut." *National Tax Journal* 61(1): 105–127.