NBER WORKING PAPER SERIES

ESTIMATING THE PRESENT VALUE OF R&D TAX BENEFITS IN THE UNITED STATES

Brandon Pecoraro
Nicholas C. Hoffman
Martin Lopez-Daneri
Elena C. Derby
Rachel Moore
Shannon E. Sledz

Working Paper 35208
http://www.nber.org/papers/w35208

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2026

This research embodies work undertaken for the staff of the Joint Committee on Taxation, but as members of both parties and both houses of Congress comprise the Joint Committee on Taxation, this work should not be construed to represent the position of any member of the Committee or the United States Congress. This work is integral to the work of the Joint Committee on Taxation and its ability to model and estimate the microeconomic and macroeconomic effects of tax policy changes. Author order reflects relative contribution. Brandon Pecoraro is listed first. Nicholas Hoffman and Martin Lopez-Daneri, who served as the primary contributing authors, follow in alphabetical order. The remaining authors are listed alphabetically. We thank Alan Auerbach, Thomas Barthold, Matthew Comey, Connor Dowd, Tim Dowd, Christopher Giosa, Benjamin Jones, Paul Landefeld, Joshua Lerner, Jacob Mortenson, David Splinter, Tracy Watkins, and Heidi Williams for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Estimating the Present Value of R&D Tax Benefits in the United States
Brandon Pecoraro, Nicholas C. Hoffman, Martin Lopez-Daneri, Elena C. Derby, Rachel Moore, and Shannon E. Sledz
NBER Working Paper No. 35208
May 2026
JEL No. D22, H25, O30, O38

ABSTRACT

Using a panel of confidential corporate tax returns, we provide the first direct estimates of the realized present value of corporate tax benefits from R&D credits and deductions in the United States. Realized tax benefits can deviate from statutory tax benefits because firms in loss status are typically unable to fully utilize credits and deductions to offset current-year taxes and instead must carry these attributes forward. We develop a novel procedure to track the intertemporal firm-level utilization of tax attributes generated by corporate R&D spending, and find that the present value of R&D tax benefits varies substantially with firms' loss status, age, and size. Old and large firms typically use R&D tax benefits quickly, while young firms – especially those that are small – frequently operate in loss status and use tax attributes more slowly. From 2012–2016, the average firm generated $0.41 in statutory tax benefits per dollar of R&D investment, with a realized present value of $0.36. Young and small firms in a loss position realized only $0.23 per dollar, a 44% decrease relative to the statutory benchmark.

Brandon Pecoraro
Joint Committee on Taxation
brandon.pecoraro@jct.gov

Nicholas C. Hoffman
Joint Committee on Taxation
Nicholas.Hoffman@jct.gov

Martin Lopez-Daneri
Joint Committee on Taxation
mlopezdaneri@gmail.com

Elena C. Derby
Joint Committee on Taxation
elena.derby@jct.gov

Rachel Moore
Joint Committee on Taxation
Rachel.Moore@JCT.gov

Shannon E. Sledz
Joint Committee on Taxation
Shannon.sledz@jct.gov

1 Introduction

The United States is among the most innovative economies in the world: it is the largest spender on research and development ("R&D") — $956 billion in 2023 in Purchasing Power Parity terms — and remains the global leader in high-value, internationally-oriented patenting activity.[1] As in other advanced economies, private R&D in the United States is subsidized through tax incentives. Economic theory offers a clear rationale for such policies: knowledge spillovers generate positive externalities (Romer, 1990), and innovation-driven creative destruction is a key engine of long-run growth (Aghion and Howitt, 1992). Furthermore, there is strong empirical evidence that tax incentives increase both R&D spending, and innovation (Bloom et al., 2019).

Accordingly, the United States federal income tax code provides substantial support for private business R&D spending. Business R&D spending can generate an income tax deduction, which reduces taxable income, and a nonrefundable tax credit, which reduces tax liability subject to statutory limitations. These provisions are economically significant in scale: for C-corporations alone, we estimate that the R&D spending that qualifies for both the deduction and credit generated claims of at least $110 billion in 2016 — about four times larger than direct federal funding for business R&D.[2]

Despite the scale and importance of federal tax incentives for R&D spending in the United States, little is known about how the statutory parameters chosen by policymakers translate into realized tax benefits. While profitable firms typically use the deductions and credits they generate to reduce current taxable income and tax liability, firms in a loss position often cannot. Because unused deductions and credits are generally not paid out immediately, the corporate income tax system features *asymmetries*: the realized value of tax benefits depends on whether firms are profitable or in loss status, and more broadly on their past, present, and expected future tax positions.

Measuring the gap between statutory and realized tax benefits is critical for assessing the incentive and distributional effects of corporate R&D tax provisions. For firms, one key component of the cost of marginal R&D capital is the *net acquisition price*: the gross price of the investment less the expected present value of tax savings generated by the outlay itself. If statutory provisions do not map one-for-one into realized tax benefits, then a measure of cost constructed solely from statutory parameters will misstate the effective tax subsidy to R&D investment. Moreover, the divergence between statutory and realized tax benefits generated by R&D spending may vary systematically across firm states — such as loss status, age, and size — shaping the distribution of incentives across firms themselves.

In this paper, we provide what to our knowledge is the first direct firm-level estimates of the realized present value of corporate tax benefits from R&D credits and deductions using confidential Internal Revenue Service ("IRS") tax return data.[3] To estimate this *R&D tax shield* across firms,

[1] See OECD Main Science and Technology Indicators and OECD Science, Technology and Innovation Outlook.

[2] See National Center For Science and Engineering Statistics, National Science Foundation.

[3] We focus on C-corporations because they account for over 90% of R&D credits generated in a typical year, which is estimated from the 2012-2016 samples of corporate income tax returns constructed by the Statistics of Income division of the Internal Revenue Service.

we develop and implement a novel approach that measures the observed generation and utilization of R&D-related deductions and credits for each annual cohort from 2012 through 2016 — including estimation of the intertemporal utilization of associated tax attributes carried forward. Critically, our estimates account for differences in carryforward *vintage*: the tax code requires that attributes carried forward must be used in the order in which they were generated. This rule has meaningful incentive effects: for firms with many years of losses, the accumulation of old-vintage carryforwards pushes the realization of any new-vintage benefits into the far future, further reducing the present-value tax benefit of a current dollar of R&D expenditure. While our realized tax-shield estimates are ex post rather than ex ante measures of investment incentives, they are informative about the utilization frictions that drive wedges between statutory tax parameters and firms' expected realized tax benefits.

Previous work has studied how corporate tax asymmetries drive a wedge between statutory R&D tax parameters and effective incentives,[4] whether through simulated tax status (Altshuler, 1988) or reduced-form user-cost adjustments for delayed utilization (Rao, 2016).[5] In contrast, by going further and directly estimating usage at the firm level we are able to uncover three central facts. First, young firms that conduct R&D spending are significantly more likely to operate in tax-loss status over multiple years. Second, young and small firms exhibit particularly high R&D intensity, with R&D spending comprising a relatively large share of total deductible expenses. Third, these same firms face the largest realization wedges: loss status pushes R&D deductions and credits into carryforwards that are used only gradually (if at all), materially reducing the present-value tax shield and increasing the realized net acquisition price of R&D for these firms.

Considering the R&D deduction and credit jointly over 2012–2016, we estimate that the average firm generates about \$0.41 in current statutory tax benefits per dollar of R&D investment; the present value of these tax benefits is about \$0.36 per dollar. This value varies substantially by firm age, size, and loss status: the present-value R&D tax shield falls to about \$0.27 on average for firms in loss status during the year of the investment outlay, and to roughly \$0.23 on average for young and small firms in loss status. While firms realize about 97% of the tax value of the R&D deduction in the year it is claimed, they contemporaneously realize only about 56% of tax value of R&D credit claims. These patterns imply that the effective tax subsidy differs systematically across firm states and is concentrated among profitable firms with sufficient current tax liability to utilize deductions and credits near the time they are generated. In addition, because a substantial share of credit dollars is realized only through carryforward usage, our estimated utilization paths imply that the revenue effects of changes to R&D tax provisions can have substantial mechanical timing "tails."

Our results carry implications for the literature on credit refundability. Returning unused credits as refunds to firms with low or zero tax liability mechanically raises the present value of the R&D

[4]Seminal work on corporate tax asymmetries includes Auerbach (1986), Auerbach and Poterba (1987), and Altshuler and Auerbach (1990).

[5]Goodman et al. (2023) measure realized firm-level net-operating-loss tax asymmetries for S-corporations, while Zwick (2021) measures firm-level realized present-value wedges implied by the choice of C-corporations to forgo net-operating-loss carrybacks.

tax shield by accelerating utilization that would otherwise occur only through delayed carryforward use, if at all. The firms for which this change matters most on the incentive margin are young and small firms, because they are most likely to face persistent low-utilization states. However, the firms accounting for most credit dollars not used contemporaneously are larger firms in loss positions, so a broad refundability policy can have most of its static fiscal cost accrue to those firms. Our estimates therefore highlight a distinction between the firms experiencing the largest change in the effective tax subsidy and the firms accounting for most of the revenue cost of refundability. Furthermore, if spillovers are larger among larger, more connected firms, policymakers face a tradeoff between relieving utilization frictions and targeting high-spillover innovators (Bloom et al., 2019).

The remainder of the paper proceeds as follows. Section 2 describes the mechanics of R&D tax provisions under tax asymmetries. Section 3 introduces the administrative data and measurement framework. Section 4 documents firm-level heterogeneity in utilization and realized tax shields. Section 5 analyzes implications for user-cost measurement and revenue effects. Section 6 concludes.

2 Mechanics of R&D Tax Provisions With Tax Asymmetries

To understand why statutory and realized R&D tax benefits can diverge, it is necessary to examine the interaction between R&D provisions and the broader corporate income tax system. The federal tax code provides incentives for business R&D spending through an income tax deduction and a nonrefundable tax credit. Firms are permitted to deduct general R&D expenses from their taxable income, and can claim a credit based on certain "qualified" R&D spending.[6] Figure 1 summarizes the ordering of relevant computations: deductions reduce taxable income and therefore income tax liability, while credits are subsequently applied to reduce income tax liability directly.

The extent to which these provisions ultimately reduce current-year tax liability depends on the interaction with limitations governed by the Net Operating Loss ("NOL") and General Business Credit ("GBC") rules.[7] An NOL arises when deductions exceed income, and the excess gets carried across years as a deduction that reduces taxable income in other periods, subject to statutory limitations. The GBC regime plays an analogous role for business tax credits, governing both how much credit may be used in a given year and how unused credits may be carried backward or forward. Tax asymmetries arise because these rules cause both the current value and present discounted value of a deduction or credit to depend on a firm's tax position: while firms with sufficient taxable income can typically monetize tax benefits immediately, firms in loss status or subject to credit limitations must often defer realization through carryforwards. We discuss these rules in Section 2.1 and illustrate their implications with a simple two-firm, two-year example in Section 2.2.

[6] See Appendix A.1 for further details about the statutory schedule that determines the value of the R&D credit.

[7] For 2012–2016, the R&D deduction and credit were provided under §174 and §41 of the Internal Revenue Code, while NOL and GBC rules are specified in §172 and §38. In addition, §280c limits the maximum combined tax benefit for a corporation that claims both the deduction and the credit on the same set of R&D expenses.

Figure 1: Mechanics of Deduction and Credit Usage with NOL and GBC Constraints

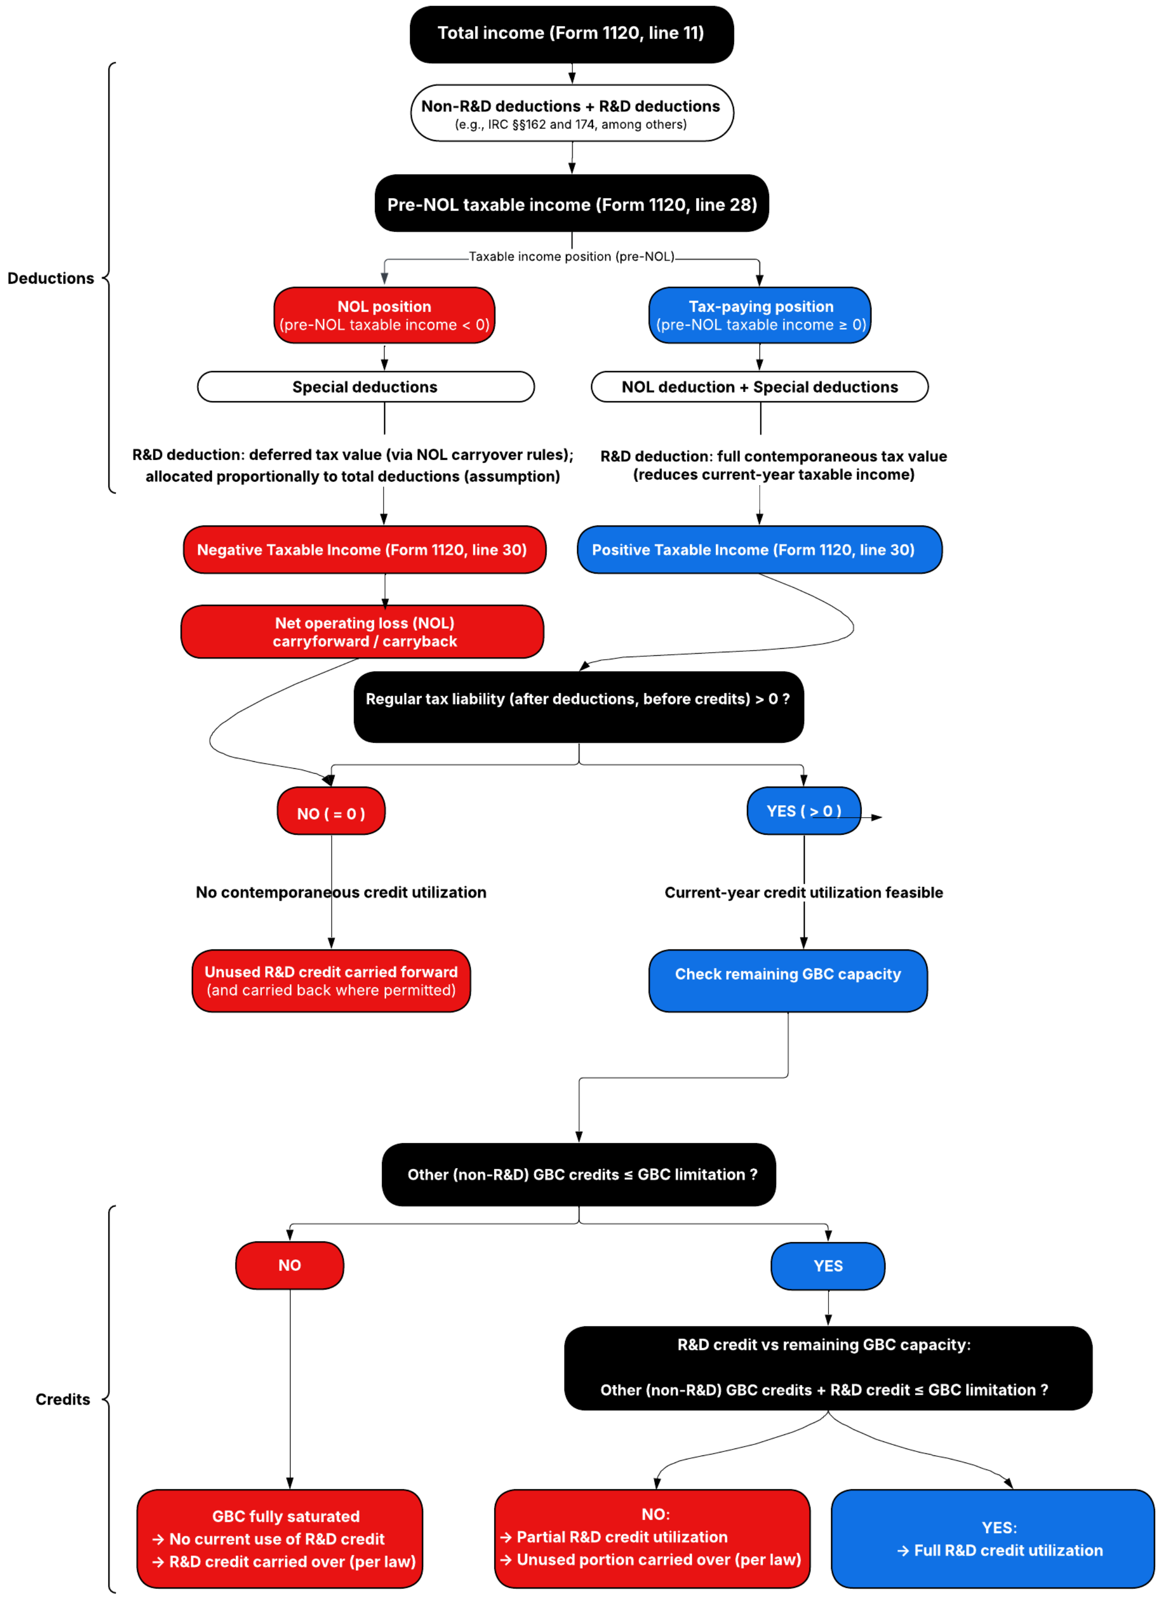

2.1 Net Operating Loss and General Business Credit Rules

A corporation's ability to contemporaneously utilize all deductions is determined primarily by whether it has positive taxable income before any NOL deduction ("pre-NOL taxable income"). When pre-NOL taxable income is positive, the firm is in non-loss status and can fully utilize all deductions generated in the current year (excluding NOL and special deductions). When pre-NOL taxable income is negative, the firm is in loss status and some or all of its current-year deductions may not generate an immediate tax benefit. Under a proportional attribution assumption, the portion of the loss attributable to the R&D deduction becomes an NOL carryover, which can be used to offset income in other years.[8] Prior to 2018, firms could choose to carry losses *back* for up to two years, and carry any remaining amount *forward* to be used as an NOL deduction for up to 20 years before the tax attributes expire. We refer to first-year deduction utilization as "incomplete" when any portion of the current-year deductions is not reflected in a current-year reduction in taxable income or used as an NOL carryback; instead such deductions become an NOL carryforward.

Central to our framework is adherence to the First-In, First-Out ("FIFO") rule:[9] any stock of preexisting "old-vintage" NOL carryforwards (accumulated from prior-year losses) must be exhausted before "new-vintage" NOL carryforwards can be applied. We track utilization of each vintage of NOL carryforwards separately, and attribute only utilization of new-vintage carryforwards to the present-value tax shield generated by a given firm-year R&D outlay. Holding the future path of pre-NOL taxable income fixed, a relatively large old-vintage stock of NOL carryforwards delays utilization of new-vintage carryforwards and reduces the present-value tax shield. Under FIFO accounting, the presence of large old-vintage carryforwards can effectively reduce the present value of tax benefits generated by new R&D spending.

A corporation's ability to contemporaneously utilize the R&D credit depends on whether it is constrained by the GBC limitation, which restricts the total amount of GBC usage — inclusive of carryforwards, current-year credits, and carrybacks —to approximately 75% of the firm's pre-GBC net income tax liability.[10] When unconstrained, a corporation can fully utilize its current-year R&D credit. When constrained, the amount of unused current-year GBCs becomes a carryover: a GBC carryover must first be carried back one year with any remaining amount carried forward for up to 20 years before the tax attributes expire. We refer to first-year credit utilization as "incomplete" when any portion of the current-year R&D credit is carried forward.

Two additional rules link deduction and credit realization. First, any old-vintage stock of NOL carryforwards must be exhausted before any current-year or carryforward GBCs may be used; NOL deductions that eliminate a firm's taxable income mechanically reduce income tax liability to zero, causing the GBC limitation to bind. Second, the GBC regime mandates a separate FIFO rule such

[8] The Internal Revenue Code does not specify an application order across individual deductions. We therefore assume that the R&D portion of deductions is proportional to its share of total deductions.

[9] Treasury Regulation §1.172-4(b)(1) specifies that NOL deductions are applied in the order of taxable years from which such losses are carried forward or back, beginning with the loss for the earliest taxable year.

[10] For our sample period, total allowable "non-specified" GBCs are generally limited to the excess of net income tax over the greater of the tentative minimum tax or 25% of the excess (if any) of net regular tax liability over $25,000. Net income tax is the sum of regular tax liability and alternative minimum tax liability less non-GBC credits.

that credits carried forward from earlier years must be used before credits generated in the current or later years; old-vintage GBC carryforwards must be completely exhausted before new-vintage credits can be applied.[11] Considering these rules jointly, firms with large stocks of old-vintage NOL and GBC carryforwards can face substantial queuing that delays the realization of tax benefits generated by R&D spending. These delays lower the present-value R&D tax shield and increase the net acquisition price relative to what is implied by statutory parameters. The same amount of R&D spending can therefore generate materially different realized tax benefits across firms.

2.2 Illustrative "Two-Firm, Two-Year" Example

Year One: Generation of New-Vintage Tax Attributes. Consider two corporations with the following identical *current-year* characteristics in year one: total income of \$100 and total deductions of \$125, the latter of which contains \$25 of R&D expenses so that no-R&D expenses are \$100. Pre-NOL taxable income is $\$100 - \$125 = -\$25$ implying that both firms are in loss status and have no regular income tax liability. Assuming that both firms elect to forgo the NOL carryback, no portion of the loss generates a first-year tax benefit through prior-year offset, and the total loss of \$25 becomes a new-vintage NOL carryforward. This is summarized in Table 1.

In this scenario, R&D expenses represent $(\$25/\$125) \times 100 = 20\%$ of total deductible expenses in year one. Under our proportional attribution assumption, the \$100 of deductions absorbed by current-year income are therefore allocated as \$80 of non-R&D deductions and \$20 of R&D deductions. The remaining \$25 loss carried forward as a new-vintage NOL is correspondingly allocated as $\$100 - \$80 = \$20$ of non-R&D deductions and $\$25 - \$20 = \$5$ attributable to R&D.

Table 1: Year One: Generation of New-Vintage NOL and R&D Credit (Common to Both Firms)

Year One Item	**Value**
Total income	100
Total deductions (incl. R&D deduction of 25)	125
Taxable income pre-NOL	−25
Tax before credits	0
New NOL carryforward created (total)	25
New NOL carryforward attributable to R&D	5
R&D credit generated (6% of R&D spending)	1.50
R&D credit carried forward to year 2	1.50
R&D deduction used in year 1	20
R&D credit used in year 1	0

Notes: The share of deductions attributable to R&D is computed as 25/125; multiplying this share by the year-one loss of 25 yields a new NOL carryforward attributable to R&D of 5. The effective R&D credit rate is assumed to be 6%, so the year-1 R&D credit generated equals $0.06 \times 25 = 1.50$. Because pre-NOL taxable income is negative, regular tax liability is zero in year one and no credits can be used.

[11] GBCs generated within a given year are applied in a particular order, with the R&D credit being applied second.

Suppose that for both firms the effective R&D credit rate is 6% so that \$25 of R&D expenses generate an allowable R&D credit in the amount of \$1.50.[12] Because year-one regular tax liability is zero, none of the R&D credit can be used in the current year. Assuming neither firm has tax liability in the prior year, the full \$1.50 R&D credit becomes a new-vintage GBC carryforward.

Year Two: Utilization and New- vs. Old-Vintage Queuing. Suppose that both firms have the same pre-NOL taxable income of \$50 in year two, but differ in the size of their *old-vintage* carryforward stocks entering year two (which were pre-existing in year one).[13] Firm 1 enters year two with old-vintage NOL and GBC carryforward stocks of \$10 and \$2 respectively. Firm 2 enters year two with old-vintage NOL and GBC carryforward stocks of \$60 and \$20 respectively.

Under FIFO rules, Firm 1's \$10 of old-vintage NOL carryforwards are applied first against its \$50 of pre-NOL taxable income, leaving \$40 of taxable income. Firm 1 then uses its new-vintage NOL carryforward of \$25, including the \$5 R&D portion, leaving taxable income of \$15. Under a statutory corporate income tax rate of 35%, pre-credit tax liability is $0.35 \times \$15 = \$5.25$. With a simplified GBC limitation equal to 75% of pre-credit tax liability, $0.75 \times 5.25 = 3.94$, Firm 1 first uses its \$2 old-vintage GBC carryforward. The remaining credit capacity is $3.94 - 2 = 1.94$, which is sufficient to fully utilize its \$1.50 new-vintage GBC carryforward attributed to R&D. Firm 1 will have no remaining old- or new-vintage carryforwards when it enters year three.

Under FIFO rules, Firm 2's \$60 of old-vintage NOL carryforwards are applied first and offset its \$50 of pre-NOL taxable income. Firm 2 cannot use any of the new-vintage NOL carryforwards, including the R&D portion, which will be carried forward to year three along with the remaining \$10 of old-vintage NOL carryforwards. With zero taxable income after NOLs, Firm 2 has no tax liability and therefore cannot use any of its old-vintage GBC carryforwards or its new-vintage R&D-related GBC carryforward; all are carried forward to year three. This is summarized in Table 2.

Present-Value R&D Tax Shield: While the statutory value of R&D tax benefits generated by both firms in year one is \$10.25 ($0.35 \times \$25 + \$1.50$), realized benefits are smaller than the statutory amount because neither firm contemporaneously utilizes all benefits in the first year. Since Firm 1 utilizes all new-vintage carryforwards by the end of year two, its realized R&D tax shield in this example can be computed directly from year-one and year-two usage. For Firm 1, let year-one and year-two usage be $use_{1,1} = 0.35 \times \$20$ and $use_{1,2} = 0.35 \times \$5 + \$1.50$ respectively, and future usage is discounted by an assumed factor of $\beta = 0.96$. Firm 1's present-value tax shield is then:

$$use_{1,1} + \beta use_{1,2} = 0.35 \times \$20 + 0.96\,(0.35 \times \$5 + \$1.50) \approx \$10.12$$

which is about 1.3% smaller than the statutory value of its benefits. On a per-dollar basis, the statutory and effective tax benefits are about \$0.41 and \$0.40 respectively, which implies that the

[12]For firms observed in our panel over 2012–2016, the sample-weighted average effective credit rate is about 6.5%. See Appendix A.1 for a description of how statutory credit amounts are determined.

[13]Note that the presence of the old-vintage carryforwards does not affect year 1 outcomes in this example because both firms have zero regular tax liability in year one.

realized net acquisition price to Firm 1 was about 0.9% higher due to tax asymmetries.[14] Since Firm 2 carries all new-vintage R&D carryforwards into year 3, its realized present-value tax shield will be strictly smaller than Firm 1's while its realized net acquisition price wedge will be larger.

Table 2: Year Two: Old-Vintage Delay of New-Vintage Utilization

	Firm 1	**Firm 2**
Tax Attributes entering year 2		
Old-vintage NOL CF stock	10	60
Old-vintage GBC CF stock	2	20
New-vintage NOL CF from year 1 (total)	25	25
New-vintage NOL CF attributable to R&D	5	5
New-vintage R&D credit CF	1.50	1.50
NOL utilization		
Taxable income pre-NOL	50	50
Old-vintage NOL CF usage	10	50
Remaining taxable income after old-vintage NOL	40	0
New-vintage NOL CF usage (total)	25	0
New-vintage NOL CF usage attributable to R&D	5	0
New-vintage NOL CF attributable to R&D remaining	0	5
Taxable income after all NOL	15	0
Tax before credits (35% rate)	5.25	0
GBC utilization		
GBC limitation (75% of tax before credits)	3.94	0
Old GBC used	2	0
Residual cap available for new credits	1.94	0
New-vintage GBC CF usage attributed to R&D	1.50	0
New-vintage GBC CF attributed to R&D remaining	0	1.50

Notes: Both firms have pre-NOL taxable income of $50 in year two and enter with same new-vintage carryforwards ("CF"). Firm 1 enters with positive old-vintage NOL and GBC CF, but old-vintage NOL stock is not large enough to eliminate taxable income; Firm 1 retains positive pre-credit tax liability and sufficient GBC limitation to utilize both its old-vintage and new-vintage GBC CF in year two. Firm 2 uses old-vintage NOL deductions to offset its year-two taxable income, leaving zero tax liability and credit capacity; neither old-vintage nor new-vintage GBCs can be used in year two, and the new-vintage R&D -related stocks are carried forward to year three.

[14]This is computed as $\frac{1-(10.12/25)}{1-(10.25/25)} \approx 1.0088$.

3 Data and Measurement

We use confidential administrative data from the Statistics of Income ("SOI") corporate stratified sample.[15] Firms in the SOI sample are weighted using inverse-probability weights so that the repeated cross-section is representative of the universe of C-corporations in the United States. From this repeated cross-section, we construct a balanced panel of corporations by identifying the firms that appear in all years 2010 through 2017. Since our goal is to track the subsequent utilization of tax attributes generated by annual R&D spending, requiring firms to appear in all panel years imposes a common sample-selection rule for every potential investment cohort and ensures that differences in observed utilization across cohorts reflect the panel window rather than firm entry or exit. We end the panel in 2017 in order to exclude any effects from the major corporate tax changes enacted under the P.L. 115-97, which would otherwise introduce confounding changes to the treatment of R&D expenditures.[16] The resulting panel includes over 22,000 tax filers representing nearly 650,000 weighted tax filers in a typical year.

The unit of observation is a firm–cohort pair in which a corporation's annual R&D spending generates R&D deductions and/or credits. There are about 15,000 unweighted firm–cohort observations in our panel (about 27,000 by weight). For each observation, we refer to "year one" as the year in which the R&D spending occurs. As described in Appendix A.1, we identify current R&D spending based on the expenditures reported for purposes of the credit because the R&D deduction is not directly observable. We then track the generation and subsequent utilization of the deductions, credits, and resulting NOL and GBC carryforwards associated with that year-one R&D spending. In order to identify new-vintage carryforwards separately from any carrybacks, we need two pre-year-one years and one post-year-one year for every firm–cohort observation. With this structure, each firm appearing in our panel may generate up to five observations — one cohort for each of the years 2012–2016 — depending on how many years we observe the firm with positive R&D expenditures. The 2012 cohort of observations has the most post-year-one years observed in the panel before it is right-censored, whereas later cohorts have fewer post-year-one years. Observed usage is calculated from the panel in a similar fashion to the illustrative example in Section 2.2; the procedure is detailed in Appendix A.

Due to selection into the SOI sample and our requirement that firms appear in all panel years, firms in our balanced panel tend to be large in terms of assets and disproportionately profitable. However, because the majority of R&D credits and deductions are claimed by large C-corporations, our panel covers large shares of the relevant aggregates in each year. For 2012–2016, our panel consistently captures about 40% of C-corporate filers by weight, and — because the sample skews toward larger corporations — about 70% of aggregate assets and 65% of gross receipts. Most importantly, over 2012–2016 our panel accounts for roughly 70–72% of total credit-eligible R&D

[15] Internal Revenue Service, Statistics of Income Division, corporate income tax return data, Tax Years 2010—2017.

[16] Although the R&D credit was enacted in 1981 as a temporary provision, Congress repeatedly renewed it — often retroactively — before making it permanent in the PATH Act of 2015. By the early 2010s, repeated extension made continued renewal a plausible expectation, even as the staff of the Joint Committee on Taxation emphasized that temporary status still created policy uncertainty (Hatch, 2002; Joint Committee on Taxation, 2011, 2015).

expenditures from the universe of U.S. C-corporations. The remaining 28–30% of R&D expenditures are generated by firms that are less likely to appear in the SOI stratified sample in every year and are therefore underrepresented in our balanced panel. To the extent these omitted R&D dollars are disproportionately generated by small and young firms — as suggested by the SOI sampling probabilities and the concentration of R&D spending among large filers — our utilization estimates may be biased upward relative to the full population, since small and young firms tend to utilize tax attributes slowly. We discuss generalizability of our results in Section 5.3.

3.1 Addressing Right-Censoring

It is not uncommon for an observation in any given year-one cohort to have tax attributes that have not been fully utilized by the time the panel is right-censored. For each such observation, we must extrapolate carryforward usage until the tax attributes are fully utilized or expire. In the case of carryforwards that remain unused when the sample ends in 2017, we extrapolate usage at the firm level, simulating usage post-sample until the carryforwards are either exhausted or expire after 20 years. In such cases, an observation's R&D tax shield is computed from a mixture of observed and extrapolated tax attribute usage. This process implies that earlier investment cohorts have more observed years than later cohorts. For example, we observe six years of usage for any carryforwards generated in 2012, while we only observe two years for those generated in 2016.

For purposes of extrapolating usage of any remaining old-vintage NOL (GBC) carryforwards, we bin firms by the number of years for which they are observed in loss status (limited by GBC constraints) over 2012–2016, separating firms with zero, one, two, and three-or-more loss (limited) years. We compute average annual usage rates within each bin, and assume that firms with unused old-vintage carryforwards at the end of the sample period continue to utilize such carryforwards at their bin-specific rate, which then converges gradually toward the sample mean. This assumption is intended as a conservative reduced-form extrapolation: it preserves persistent differences across firms with different histories of loss or credit limitation, while avoiding the polar assumptions that end-of-sample status either persists permanently or disappears immediately. Specifically, the bin-specific usage rates initially assigned to firms at the end of the sample period converge toward the sample mean, so that firms observed in loss status throughout the sample are not mechanically assumed to remain in loss status forever. This process strikes a balance: it preserves heterogeneity across firms with different histories of losses while avoiding the restrictive assumption that firms observed in loss status throughout the sample will *never* leave loss status. Throughout the extrapolation process we maintain strict enforcement of FIFO, allowing utilization of new-vintage carryforwards only once old-vintage carryforwards are deemed exhausted. The full procedure is detailed in Appendix B.

3.2 Weighting

Since our balanced panel is constructed by selecting firms that appear in all years 2010 through 2017 of the SOI repeated cross-section, the interpretation of any weighted statistic depends on whether it is intended to describe (*i*) a characteristic of typical surviving firm appearing in the panel, or (*ii*) an aggregate quantity. For measures that show dispersion of *firm-level* characteristics, we use only SOI sample weights.[17] For measures that describe *aggregate* usage of deduction, credit, or tax attribute quantities, we use SOI sample weights multiplied by the dollar quantity of interest (i.e., dollar-sample weights) so that the aggregates are representative of dollar totals. These weights do not aggregate to the full universe of entrants and exits; rather, they describe the always-observed population used to form the balanced panel.

4 Results

We begin by documenting the aggregate generation of R&D deductions and credits. Figure 2 shows that, for the full set of domestic C-corporations in the cross-sectional SOI samples, the annual tax values of R&D deductions and credits generated over 2012–2016 averaged about $79 and $13 billion (in nominal terms), respectively.

Figure 2: R&D Credits and Deductions Generated

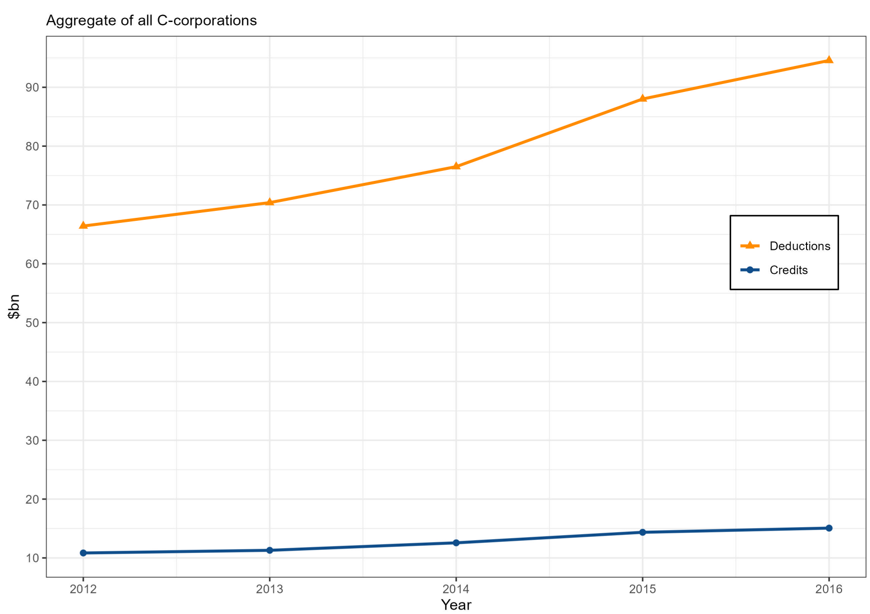

Note: All values are in nominal terms and are computed from the full set of domestic C-corporations in the SOI samples. The R&D deductions depicted here are those associated with outlays reported by C-corporations for purposes of the §41 credit, and are therefore only a subset of total §174 deductions. (See Appendix A.1 for more details.) The tax value of such deductions as reported above is estimated by multiplying the amount of deductions by 35%, which was the modal statutory corporate income tax rate over our sample period.

[17] In particular, single-year measures are weighted by the year-specific weights, while present-value measures are weighted by the corresponding year-one weights.

A main point of this paper is that contemporaneous utilization of these deductions and credits is incomplete. For the subset of C-corporations included in our panel, we estimate that about 97% of deductions and only 56% of credits generated in 2012–2016 were used to immediately reduce current-year taxable income and tax liability — top-line figures which mask substantial heterogeneity across firms. Before describing this heterogeneity in Section 4.2, we describe the relationship between firms' loss status, age, and size, which together shape the utilization patterns that ultimately determine R&D tax shields.

4.1 Firm Characteristics: Losses and R&D Intensity

We document two patterns that help explain the heterogeneity in realized R&D tax shields. First, small and young firms with R&D spending are significantly more likely to operate in a tax-loss position in a given year and to spend multiple years in loss status. Second, these firms exhibit higher R&D expenses as a portion of total deductions.

Figure 3A plots the distribution of firm ages in our sample — with a median age in 2012 of 19 years — binned by number of loss-status years over the 2012–2016 period. In the left panel, the distribution of firm age becomes more concentrated at lower values as the number of loss years increases: the median firm with zero to one loss years over the sample period is over a decade older than the median firm with four to five loss years.

The right panel of Figure 3A decomposes the left panel by firm size, as measured by total assets in 2012. Firms in blue are "small" and hold assets below the median level in 2012 — about $20 million — while firms in red are "large" in that they hold above-median assets in 2012. The differences between the blue and red densities make clear that *small* firms are an important driver of the relationship between age and loss years: the distribution of ages among firms with four or five loss years skews far younger for small firms than for large firms.

Smaller firms, regardless of age, additionally have expenses tilted more heavily towards R&D. To measure the prominence of R&D among a firm's total deductible expenses, we define *R&D intensity* as a firm's R&D deduction as a share of its total deductions within a given year. Figure 3B plots the density of R&D intensity among firms in our panel, broken down by age and size with size defined using the same groups as in Figure 3A. Additionally, firms in a loss position typically have greater R&D intensity than do non-loss firms, and this pattern is particularly pronounced for younger firms.

Taken together, Figures 3A–B establish that firms that are both young and small tend to have higher R&D intensity and more frequent loss years. Our evidence complements patterns on firm dynamics documented in data on publicly corporations, where R&D intensity is higher among younger and smaller firms even as current profitability is lower (Eisfeldt et al., 2026). In our setting, this low-profitability phase has first-order policy consequences: when firms operate in loss status, R&D deductions and credits are pushed into NOL and GBC carryforwards, and FIFO ordering and statutory limitations delay utilization and reduce the realized present value of the associated tax shields.

Figure 3: Firm Characteristics: 2012–2016.

A. Densities of Years Spent in Loss Status by Firm Age

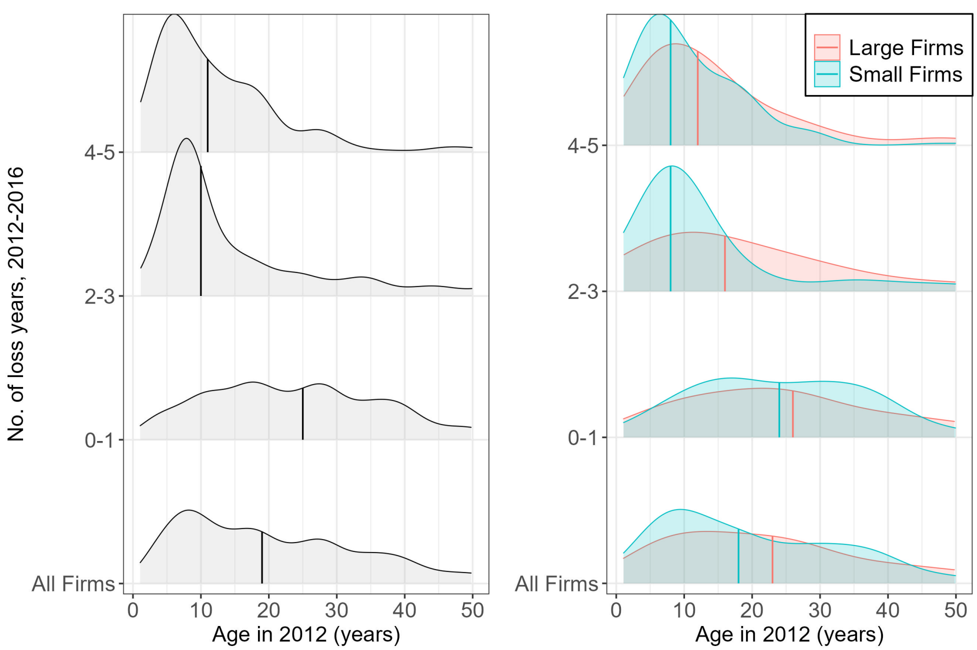

B. R&D Intensity by Firm Size and Age

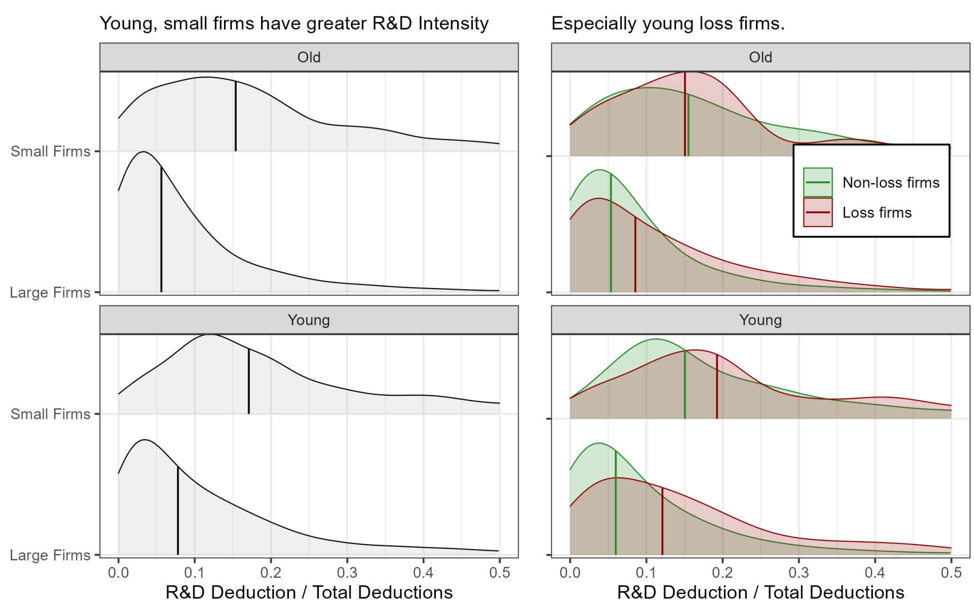

Note: Densities in Panel A are calculated using 2012 sample weights; densities in Panel B use year-specific sample weights. Young firms are those below the median age of 19 years in 2012. Small firms are those with below-median ($20m) assets in 2012. Loss firms have negative pre-NOL income. Vertical lines denote group-specific median values.

4.2 Usage in the First Year

The patterns described in Section 4.1 imply that a substantial portion of R&D deductions and credits generated in a given year — especially those generated by young and small firms that frequently operate in loss status — will not be contemporaneously used, and will instead become tax attributes that are carried forward. We next quantify the extent of this incomplete first-year utilization. Unless otherwise noted, all aggregates and shares reported in this section are computed using dollar-sample weights based on the amount of R&D deductions and credits generated.

Figure 4A shows the amount of R&D deductions and credits generated in each of our panel years, and breaks down these totals into the amounts contemporaneously used to reduce taxable income and tax liability, inclusive of carrybacks.[18] For example, in 2012 the nominal tax value of R&D deductions and credits generated across all C-corporations was about $66 billion and $11 billion respectively. Relative to deductions, a substantially smaller portion of credits are contemporaneously used: about 96% of deductions and 58% credits are used in year-one. This first-year usage pattern is relatively stable over 2012–2016 even as total R&D deduction and credit generation increases in nominal terms over the same period.

A larger proportion of deductions are used in year one for two primary reasons. First, deductions are applied before credits: firms calculate taxable income by subtracting total deductions from total income, and subsequently compute income taxes before credits. Since credits are limited from reducing pre-credit tax liability by more than approximately 75%, a larger amount of deductions relative to income increases the likelihood that firms will be constrained in the use of credits. Second, in the absence of a statutory ordering rule for deductions, our proportional-attribution assumption assigns partial R&D -deduction use to all firms with R&D expenditures in loss status. As a result, we ascribe partial usage of the R&D deduction to *all* firms with R&D expenditures in a loss position by assumption.

Pooling across years, Figure 4B shows a decomposition of first-year utilization by firm age and size, revealing significant heterogeneity. For R&D deductions, firms that are both young and small are particularly restricted in contemporaneous utilization, using only about 69% of their year-one deductions on average over 2012–2016 (inclusive of carrybacks), which is substantially lower than the 95%–99% average usage rates of the other groups. Since these firms are those most frequently observed in loss status, as described in Figure 3A, they are less able to use deductions against current- or prior-year income.

For R&D credits, firms that are both young and small are similarly more restricted in contemporaneous utilization; we estimate that about 14% of their credits are used in year-one on average over 2012–2016. Although they are somewhat less restricted, we estimate that the amount of credits used in year-one among old, small firms is also below average at about 38%. As with deductions, this reflects the fact that smaller firms in our panel are more frequently observed in loss status.

[18] We estimate that of firms eligible for carrybacks — those in a loss status with positive taxable income in the previous two years — 46% elect to carry back deductions, consistent with Zwick (2021). Among firms with R&D expenses, 59% of those eligible to take carrybacks do so.

Figure 4: First-Year R&D Deduction and Credit Usage

A. Aggregate First-Year Usage, 2012–2016

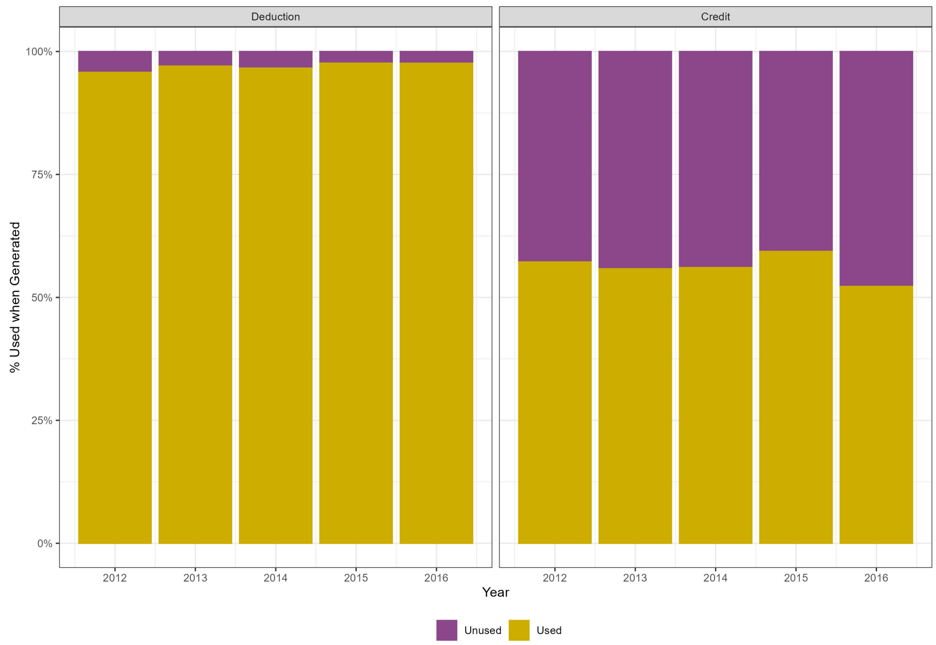

B. Distribution by Firm Characteristics

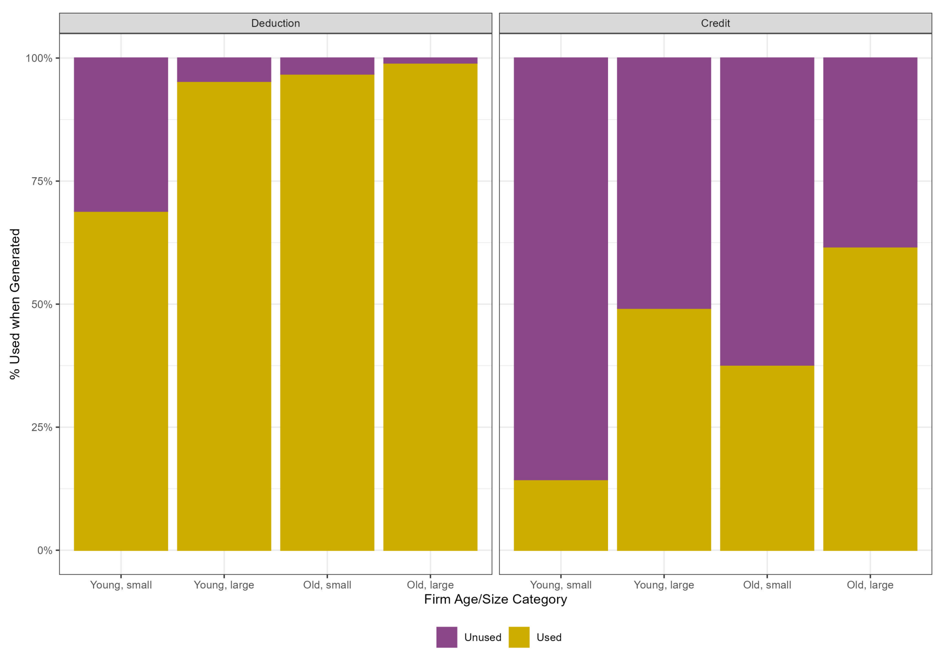

Note: First-year utilization is defined to include (i) a reduction in current-year taxable income or tax liabilities, and (ii) a reduction in prior-year taxable income or tax liabilities due to carrybacks. Small firms are those with below-median assets; young firms are those below median age. Figure B pools years 2012-2016. All values are weighted by firm-year dollar-sample weights.

4.3 Usage in Later Years

How long does it take to realize the tax benefits generated by R&D spending? To answer this question, we define *cumulative usage* of credits and deductions as a dollar-sample-weighted metric that includes both first-year usage and later-year usage of the associated tax attributes when first-year usage is incomplete.

Figure 5 shows cumulative usage in the years following the year-one R&D expenditure, again broken down by firm age and size and pooled across years 2012–2016. As a benchmark for "high" cumulative utilization, consider the subset of firms that are both large and old. These firms are least likely to be in loss status, so the vast majority of the tax benefits associated with their deductions and over 75% of those associated with their credits are used within five years of generation. Contrast this with firms that are both young and small: only about 70% and 22% of the tax benefits associated with their deductions and credits, respectively, are used within five years.

Within each group of firms, a clear gap exists between the relatively higher cumulative usage of deductions and the relatively lower cumulative usage of credits. While no group has cumulative deduction utilization below 70% even after two decades, cumulative credit usage remains in a wide range of about 30%–80%. This result highlights that even if the asymmetry between loss and non-loss positions is properly accounted for, assuming that firms immediately realize new-vintage tax benefits once they return to profitability can overstate the size of the subsidy, since firms must first use up old-vintage carryforward stocks before they can utilize new vintages. This feature of FIFO also implies that, at the firm level, credit usage is highly discontinuous: firms with GBC carryforwards typically do not use any credits for several years while exhausting their NOL carryforward stock, and then often utilize their relatively smaller new-vintage GBC stock over a short horizon. The smooth aggregate pattern of carryforward usage in Figure 5 reflects averaging across firms with discontinuous individual usage paths; for right-censored years, the extrapolation procedure is designed to preserve that firm-level extensive-intensive structure.[19]

The cumulative usage patterns in Figure 5 mix contemporaneous and delayed usage, and include firms that may have carried forward only deductions or only credits. To isolate utilization of new-vintage tax attributes that are carried forward after year one, Figures 6A and B show cumulative usage of *old-vintage* and *new-vintage* tax attributes for the subset of firms that carry forward both deduction and credit tax attributes after year one. In these figures, age is more closely associated with year-one utilization capacity (via loss status), while firm size is more closely associated with the persistence of delayed utilization through old-vintage carryforward overhang. Consequently, firms that are both young and small exhibit low initial utilization capacity and the slowest catch-up in new-vintage usage.

Here the message is clear: new-vintage usage is slow. Among firms that carry forward both deduction and credit attributes after year one — typically firms in loss status — only about 20% of new-vintage deduction carryforwards and 10% of new-vintage credit carryforwards are used within five

[19] Our extrapolation procedure used to simulate usage for years that are right-censored in the data nests a discrete extensive-intensive utilization process as described in Appendix B.2.

years. Even over the full 20-year carryforward window, only about 22% of deduction carryforwards and 20% of credit carryforwards are *ever* used before expiration (not shown). Utilization is lower still among firms that are both young and small at the time of investment: only about 6% of new-vintage deduction carryforwards and 10% of new-vintage credit carryforwards are ever used before expiration.

Figure 5: Cumulative Usage of R&D Deductions and Credits

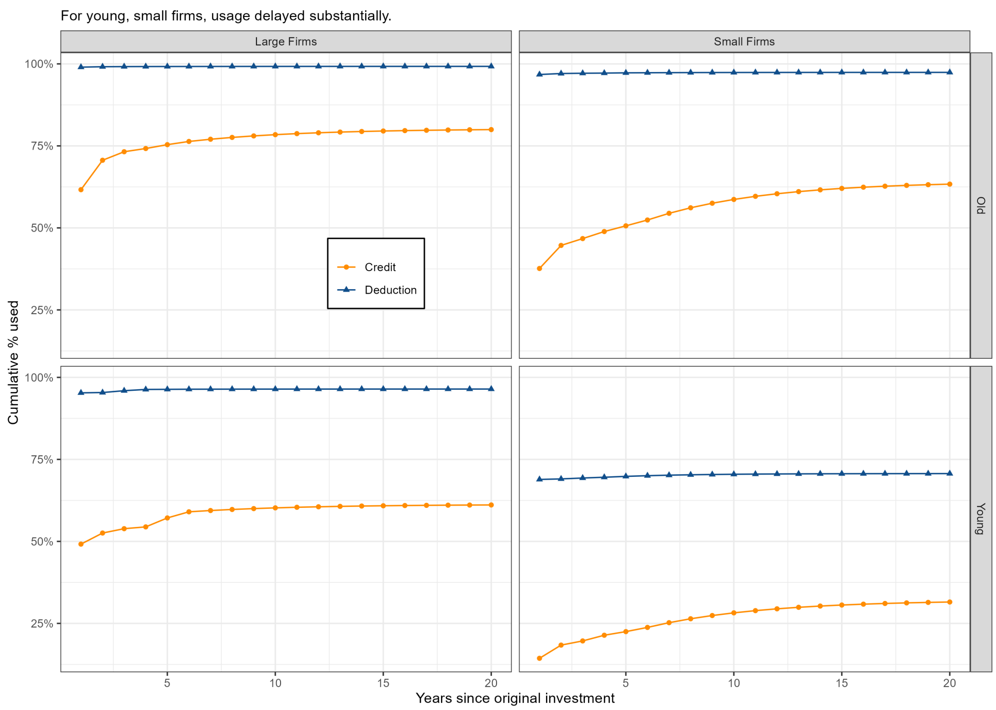

Note: Aggregate estimates are weighted by firm-year sample weights and dollars. Small firms are those with below-median assets; young firms are those below median age. Figures pool investment years 2012-2016, and extrapolate to estimate usage beyond 2016.

4.4 Present Value R&D Tax Shields

Figure 7 shows the distribution of the per-dollar present value of the R&D tax shield estimated for our panel. For each firm, we calculate this by (i) valuing the present value of deduction usage at a 35% corporate income tax rate, (ii) adding the present value of R&D credit usage, and (iii) dividing by year-one R&D expense. Present values are computed by discounting year-t utilization flows by an assumed annual factor of 0.96. The solid vertical lines in Figure 7 denote group-specific median values, while the dashed lines denote group-specific means.[20]

[20]On a per-firm basis, the maximum value that the -per-dollar measure can take depends on the firm's "base" level of qualified R&D expenditures, which varies across firms as specified by the statutory credit schedule.

Figure 6: Cumulative Usage of Carryforwards by Vintage

A. Cumulative Usage of Old-Vintage Tax Attributes Carried Forward

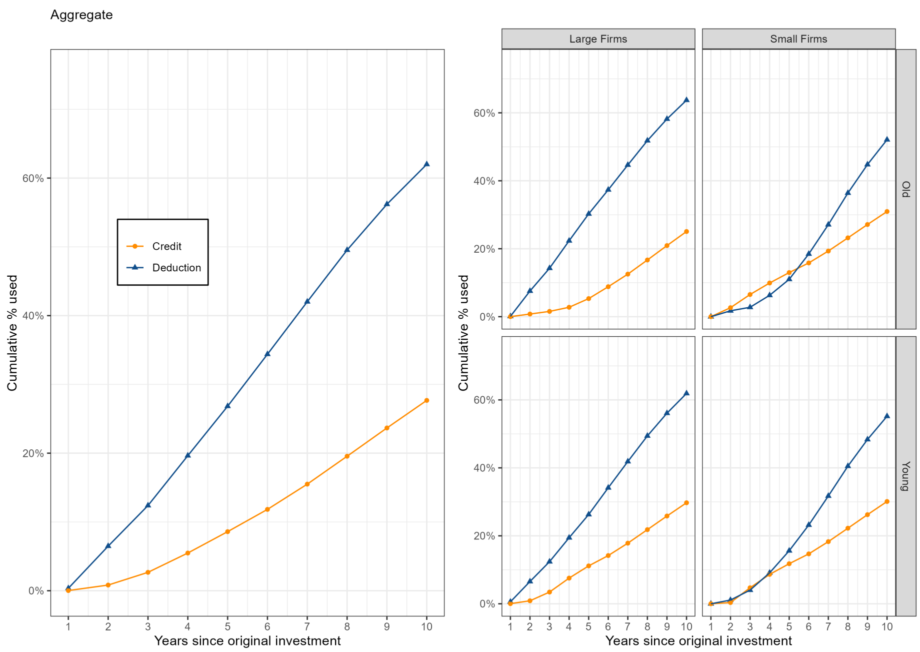

B. Cumulative Usage of New-Vintage R&D Tax Attributes Carried Forward

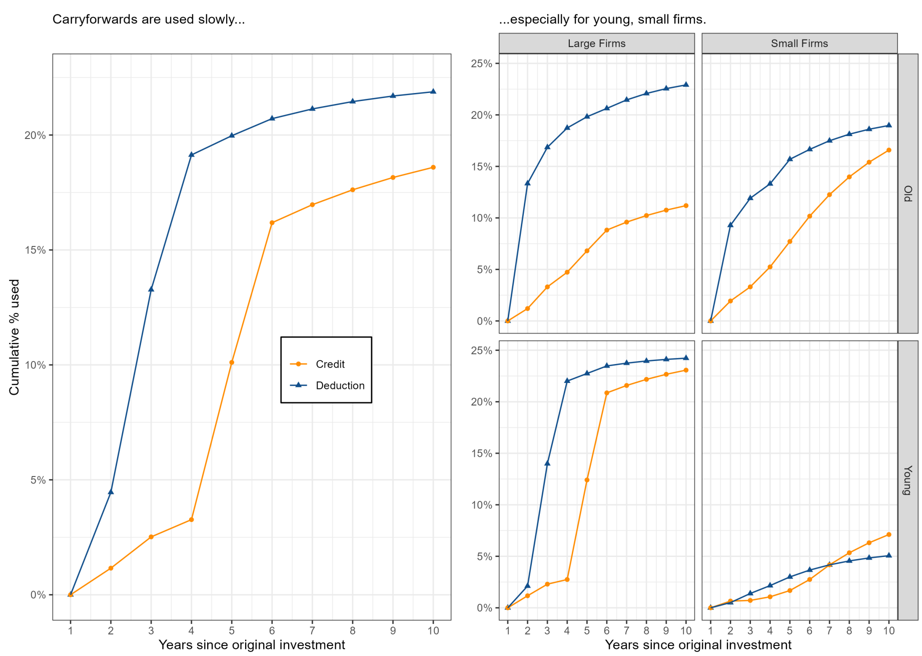

Note: All values are weighted by dollars and firm-year sample weights. Sample in Panels A and B restricted to firms carrying forward both R&D deduction *and* credit tax attributes. Small firms are those with below-median assets; young firms are those below median age. Figures pool investment years 2012-2016, and extrapolate to estimate usage beyond 2016.

Figure 7: Distribution of Present Value (PV) R&D Tax Shields

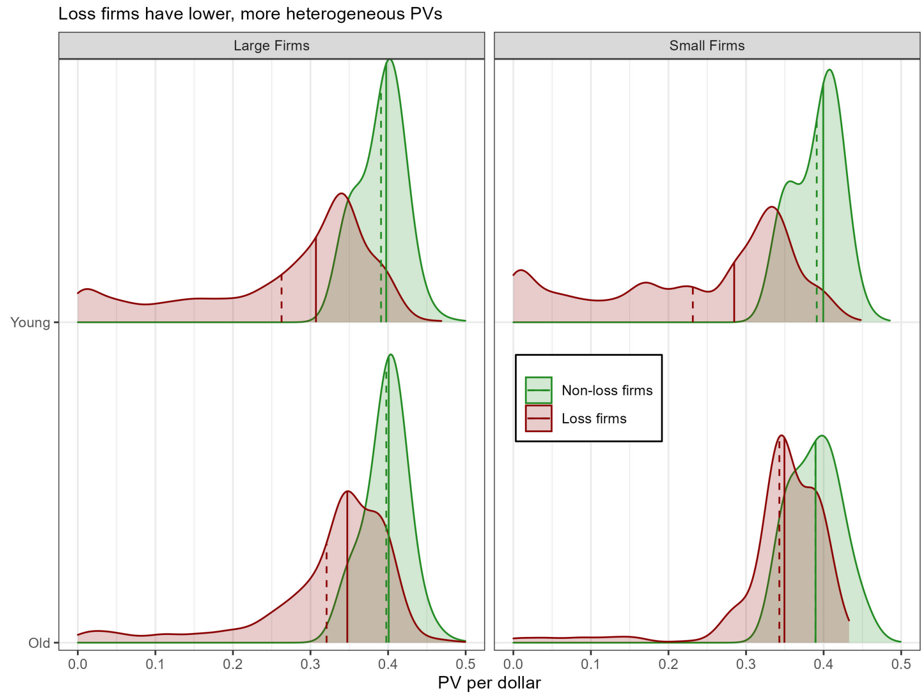

Note: Densities calculated using firm-year sample weights. Small firms are those with below-median assets; young firms are those below median age. Figures pool investment years 2012-2016, and extrapolate to estimate usage beyond 2016. Vertical lines indicate group-level medians; dashed lines indicate group-level means.

The average firm in our panel realizes \$0.36 in present-value tax benefits per dollar of R&D spending. Because operating in loss status typically causes unused credits and deductions to become carryforwards, loss-status firms of all ages and sizes tend to have lower present-value tax shields than do non-loss firms: The average non-loss-status firm realizes \$0.39 in present-value tax benefits per dollar of R&D spending, as compared to only \$0.27 for the average loss-status firm. This gap reflects the result in Figure 6B, that average utilization is slow when a deduction or credit becomes a carryforward. More generally, our finding of slow tax attribute usage for C-corporations contrasts with the relatively fast usage for S-corporations documented by Goodman et al. (2023).

The gap in present-value benefits is especially striking for the young, small firms in the top-right panel of Figure 7. The mean loss firm in this category realizes only about \$0.23 in present-value tax benefits per dollar of R&D spending, lower than the \$0.27 realized by the average loss firm. The low R&D tax shield realized by young, small loss firms — relative to all loss firms — reflects the additional frictions that these firms face in using the benefits that they generate: young, small firms often spend many years in loss status, and even once they become profitable, their accumulated old-vintage NOL and GBC carryforward stocks act as overhang, further delaying the utilization of new-vintage R&D tax benefits.

Across all firm ages and sizes, the mean present-value tax benefit per dollar of R&D expenditure is below the median for firms in a loss position, suggesting that relatively low-present-value observations pull the average down within these categories. This pattern is even more pronounced when considering the distribution of *dollars*, rather than the distribution of *firms*. When weighting by both sample weights and R&D dollars, the average dollar of R&D expenditures generates a present-value tax benefit of $0.38. Among firms in a loss position, this average per-dollar present-value tax benefit falls to $0.30, and falls further to $0.17 for young, small loss firms. These figures suggest that among young, small firms in a loss position, R&D expenditures skew towards more usage-limited firms realizing lower present-value benefits.

5 Applications

5.1 Net Acquisition Price Wedge and User Cost

In the Hall and Jorgenson (1967) user-cost framework, the present-value tax shield is the channel by which tax policy influences the marginal cost of investment.[21] However, the standard user-cost framework does not capture the effect of tax asymmetries, which can meaningfully reduce the true tax subsidy to R&D spending (Rao and Simcoe, 2025). When tax asymmetries reduce the present-value tax shield by delaying or preventing the monetization of credits and deductions, they raise user-cost through an increased net acquisition price: the price per unit of R&D capital, net of effective tax benefits generated by the marginal dollar of investment.[22]

To measure the effect of tax asymmetries on the net acquisition price of R&D capital, we define a *net acquisition price wedge* as the ratio of the realized net acquisition price to the corresponding statutory benchmark:

$$\hat{\mathcal{W}}_a = \frac{1 - \tau Z_a^{\text{real.}} - \kappa_a^{\text{real.}}}{1 - \tau Z_a^{\text{stat.}} - \kappa_a^{\text{stat.}}}$$

Letting τ denote the statutory marginal corporate income tax rate, the terms $\tau Z_a^{\text{real.}}$ and $\kappa_a^{\text{real.}}$ denote the *realized* present-value tax shield for deductions and credits, respectively, per dollar of R&D investment. Their statutory counterparts are denoted by $\tau Z_a^{\text{stat.}}$ and $\kappa_a^{\text{stat.}}$. The wedge will vary across firms because: (i) the realized tax shield depends on a firm's capacity to utilize year-one deductions and credits, as well as the speed at which it utilizes any new-vintage tax attributes carried forward, and (ii) the statutory credit value is itself endogenous to firm behavior because the credit base depends in part on past R&D spending.

[21] See Chodorow-Reich (2025) for a comprehensive review of the neoclassical investment model.

[22] The net acquisition price of capital is not the only means by which tax asymmetries affect the cost of capital. Because losses can delay tax liability into the future, asymmetries also affect the effective marginal rate on operating returns. In particular, Auerbach (1983) and Auerbach and Poterba (1987) point out that the negative effect of tax asymmetries on net acquisition price may be partially offset by the fact that these asymmetries also increase the present value of future post-tax marginal returns. Because our estimates are intended to identify only the realized tax shield generated by the R&D outlay, we do not attempt to estimate a full asymmetry-adjusted user cost.

Figure 8: Distribution of Net Acquisition Price Wedge for R&D Capital

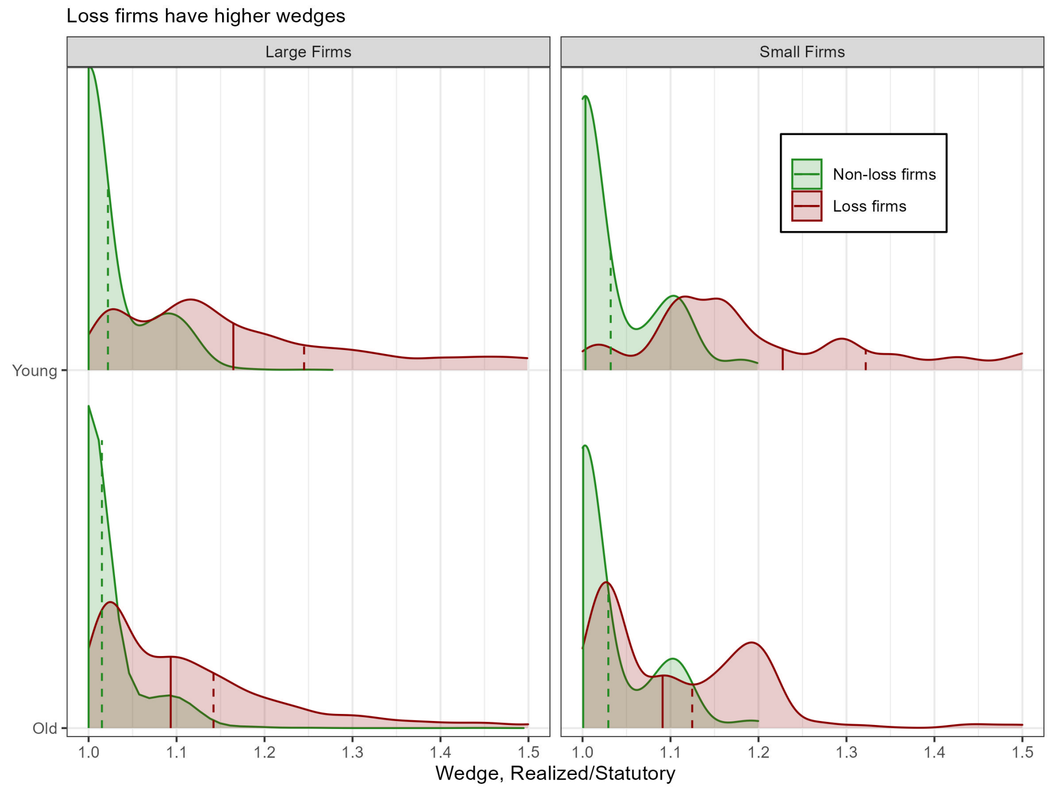

Note: Small firms are those with below-median assets; young firms are those below median age. Solid vertical lines denote group-level medians; dashed lines denote group-level means. A wedge of one indicates that realized and statutory benefits are equal. Larger wedges correspond to larger increases in effective net acquisition price due to delayed usage of tax attributes. All densities computed using firm-year sample weights.

Figure 8 shows the distribution of net acquisition price wedges for firms in our panel, again broken down by loss status, age, and size. The delayed usage of tax benefits which accompanies operating at a loss incurs a substantial increase in the post-tax price of R&D capital. Across all groups, firms that are in loss status when the year-one R&D expenditure is made (red) face a higher wedge than do non-loss firms (green). The mean loss firm, irrespective of age or size, faces a net acquisition price wedge of about 1.24. By contrast, the mean non-loss firm faces a wedge of 1.02, reflecting minimal deviation between statutory and effective cost. Once again, this discrepancy is larger for firms that are both young and small: the mean wedge for young, small firms is about 1.03 when in non-loss status, as compared to about 1.32 when in loss status. Even for large, old firms, which represent an outsize portion of aggregate R&D spending, the mean wedge is about 1.02 and 1.14 when in non-loss and loss status respectively.

The magnitude of the wedges in Figure 8 has two distinct implications. First, it implies that empirical work using statutory net acquisition prices — or user-cost series that embed statutory net acquisition prices — may suffer from *nonclassical* measurement error: such regressors omit state-dependent wedges and may therefore be correlated with both true incentives and observed R&D expenditures. Second, the wedge we construct is *ex post* — it is realized along an endogenous path of taxable income, attribute accumulation, and utilization — so it is itself endogenous to investment decisions and cannot be interpreted as the incentive firms faced *ex ante*. The appropriate ex ante incentive object for both firms' decision-making and econometric inference is therefore a forward-looking user cost that incorporates *expected* net acquisition price wedges, and more generally expected tax asymmetries, conditional on the firm's information set and state.[23]

5.2 Implications for Revenue Estimation

The observed usage patterns that we estimate have crucial implications for estimating the revenue effects of proposals to modify the tax treatment of R&D expenditures. In particular, since not all firms realize the R&D tax shield contemporaneously with deduction and credit generation, changes to these provisions will be associated with revenue effects that have a "tail" pattern. This pattern is an artifact of delayed usage rather than micro- or macro-dynamic behavioral effects, and can be illustrated with a "static" revenue estimate.[24]

Consider a year-one cohort of corporate R&D credits worth $12 billion.[25] For a baseline policy environment resembling 2012–2016 tax law, the top panel of Table 3 implies that about 56.7% (about $6.8 billion) of these credits would be used to reduce tax liability contemporaneously with generation. By the end of year ten, cumulative usage increases to about 72.2% via GBC carryforward, for a total ten-year fiscal cost of about $8.7 billion. In an alternative policy environment where unused year-one credits are refundable, the fiscal cost would be the full $12 billion of R&D credits claimed in year one. The static, ten-year revenue effect of *modifying* the R&D credit rules to allow for refundability is therefore $12 billion minus $8.7 billion — about $3.3 billion.

For purposes of this exercise, the static revenue effect corresponds to the portion of the year-one R&D credits generated that are carried forward under nonrefundability but immediately monetized under refundability. Assuming no changes to the underlying investment behavior, the benefits of refundability are distributed among the firm groups according to their share of aggregate credits that are *unused* in the year generated. Under this assumption, our estimates imply that of the $3.3 billion in fiscal costs in Table 3, about 2% accrue to young, small firms, 1% to old, small firms, and the remaining 97% to large firms of all ages.[26] Thus, while young and small firms tend to face the largest net acquisition price wedges under nonrefundability, the bulk of the static dollar revenue

[23]For estimating investment responses to the 2017 changes in the corporate tax rate and tax depreciation rules, Chodorow-Reich et al. (2024) incorporate loss-status wedges into their user-cost tax term.

[24]Note that a "static" revenue estimate differs from a conventional revenue estimate of the type produced by the staff of the Joint Committee on Taxation, which incorporates microdynamic behavior under the assumption of fixed nominal Gross National Product. For more information, see Joint Committee on Taxation (2025b).

[25]We abstract from the deduction that may also be generated by the same R&D expenditures for simplicity.

[26]Approximately 53% of the benefits of refundability accrue to old, large firms, and 44% accrue to young, large firms. Each year, large firms generate nearly 99% of all R&D deductions, and 98% of all R&D credits, with young and small firms generating approximately 1% of each. The distribution of *unused* credits and deductions skews more towards young and small firms than does the distribution of *generated* benefits because, as shown in Section 4.2, these firms are more likely to carry a given dollar of tax benefit forward.

cost of a refundability holiday accrues to large firms because they account for an outsized share of aggregate R&D credit generation.

Table 3: Illustrative Example: Partial Revenue Effect of R&D Credit Refundability

R&D Credit Usage and Carryforward Pattern *(in percent)*

Year	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Cumulative Usage	56.7	61.4	64.8	66.7	68.3	69.9	70.9	71.4	71.8	72.2
Usage Flow	56.7	4.7	3.4	1.9	1.6	1.6	1.0	0.5	0.4	0.3

Revenue Cost for Year-1 $12 billion of R&D Credit *(in millions)*

Year	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Without Refundability	6,807	564	406	224	198	192	117	60	50	41
With Refundability	12,000	–	–	–	–	–	–	–	–	–
Difference	*-5,193*	*564*	*406*	*224*	*198*	*192*	*117*	*60*	*50*	*41*
Ten-Year Total = -3,342										

Note: Cumulative usage pattern is computed as a weighted average across all year-one cohorts, smoothed with an end-point preserving three-year moving average. For averaging across year-one cohorts, weights are chosen for each cohort by the number of observable years it contributes to the panel, so cohorts with longer observed utilization paths receive greater weight. In particular, for year-one-cohorts $s = \{2012, 2013, 2014, 2015, 2016\}$ with corresponding number of observable years (including year one) $n_s = \{6, 5, 4, 3, 2\}$, we compute the weight as $\omega_s = n_s/(\sum_s n_s)$.

The fact that the majority of the fiscal costs of such a policy do *not* accrue to the firms whose net acquisition prices are most impacted by tax asymmetries captures a central tension in the refundable-credit literature. Agrawal et al. (2020) and Dechezleprêtre et al. (2023) highlight that the incentive effects of refundability are particularly salient for firms with low or zero current tax liability — often young, small firms — because it can relax near-term liquidity constraints by converting delayed carryforwards into current cash. Our estimates reflect this mechanism: refundability raises the realized present value of the R&D tax shield for low-utilization-capacity firms by accelerating utilization that would otherwise occur only through delayed carryforward use, if at all. Because firms with low utilization tend to be young and small, it is these firms that see the largest increase in per-dollar, present-value tax shield and therefore the largest decrease in net acquisition price.[27]

This tension confronts policymakers with a choice of which firms refundability policies ought to target. The welfare case for targeting young and small firms depends on the expected social returns to their R&D. For example, Bloom et al. (2013) suggest that R&D undertaken by more connected firms — which tend to be larger in size — may generate greater spillovers, implying a potential tradeoff between relieving utilization frictions and targeting high-spillover innovators.[28]

[27] An alternative to credit refundability is to allow the portion of credits not used to reduce income taxes to be used to reduce payroll taxes. The PATH Act of 2015 modified the R&D credit to allow "qualified small businesses" to apply up to $250,000 of the credit against payroll tax liability and "eligible small businesses" to claim the credit against alternative minimum tax liability. However, we observe no corporations making these elections in our panel.

[28] See Bloom et al. (2019) for a broader discussion of these design considerations and the role of financial constraints

5.3 External Validity of Results

As noted in Section 3, the SOI sampling process and our panel requirements imply that firms in our panel are likely to be older, larger, and more profitable relative to the broader universe of U.S. domestic C-Corporations. To what extent do our results apply to this broader universe?

Table 4 gives further detail on how C-corporations included in our panel compare to C-corporations excluded from our panel, as well as the broader SOI repeated cross-section sample inclusive of S-corporations. For purposes of gauging the external validity of our results, we consider only our starting year 2012 and limit firms to those with positive R&D spending in that year. The first three columns of Table 4 confirm that firms in our panel are on average older and larger than both the C-corporations that our sample selection excludes, and the average firm in the SOI corporate sample. Furthermore, C-corporations in our panel are less likely to be in a loss position than those excluded.

The comparison in Table 4 suggests that our balanced-panel estimates likely overstate aggregate utilization relative to the broader universe of C-corporations in the United States, while the direction of any bias in our distributional results is ambiguous. In terms of aggregate implications, the C-corporations that we omit are relatively smaller, younger, and more likely to be in a loss position than those we include. Furthermore, these firms' deductible expenses are tilted more heavily towards R&D spending. As such, the majority of the omitted C-corporations share characteristics with the firms in our sample for whom we infer low present values of R&D tax benefits: these firms are young and small by assets, have a high propensity to be in loss positions, and have high R&D intensity. Including these C-corporations in our analysis would likely decrease aggregate usage of R&D tax benefits. This concern is mitigated, however, by the fact that firms in our panel account for over 70% of R&D spending in each year.

The comparison nonetheless suggests that, in the broader universe of C-corporations, firms that are both young and small would continue to realize a lower present-value tax benefit per dollar of R&D expenditure. Expanding our sample to include all C-corporations present in the SOI sample in 2012, would reduce the cutoffs for "young" and "small" from 19 years to 13 years and about from $20 million to $8.6 million in assets respectively. Strikingly, about 80% of firms classified as young and small relative to these broader cutoffs are in a loss position, and the R&D deduction represents about 28% of total deductions for firms in this category. Redefining the cutoffs to include these additional C-corporations would thereby have two effects on the subgroups: some firms that are currently classified as young and small would become old and large, and a new cohort of young and small firms — who are likely to be even more limited in their usage of R&D tax benefits — would be added. The first change would reduce the gap in realized present values between young, small firms and the rest of the population, while the second would increase it. On net, although the size of the gap between young, small firms and all other firms could shift in either direction, the broader pattern would likely remain: firms that are both young and small, especially those in loss positions, would continue to realize relatively low present-value tax shields per dollar of R&D expense

among young firms.

Table 4: Comparing Panel Firms to Others in Corporate SOI - 2012

	Age	Assets	Receipts	Loss Share	R&D Exp.	R&D Intensity	Firms
	Years	*$, millions*		*%*	*$, millions*	*%*	*by weight*
Panel C-corps	24.6	4,786	1,359	35	23	15	5,680
C-corps excluded	15.0	507	204	56	5	21	8,489
All SOI C-corps	**18.9**	**2,222**	**667**	**47**	**13**	**18**	**14,169**
S-corps	23.5	30	42	18	1	17	11,119
All SOI Corps	**20.9**	**1,258**	**392**	**34**	**8**	**18**	**25,288**

Notes: Values in table are for firms in 2012 with positive R&D expenditures. All reported values are averages using firm-year sample weights. Subcategory totals may not add to aggregates due to rounding.

S-corporations differ from the C-corporations that are the focus of our analysis in ways that do not imply a clear application of our findings. As shown in Table 4, S-corporations that conduct R&D spending tend to be both smaller and older than their C-corporation counterparts, and are also less frequently in loss positions. In addition, Goodman et al. (2023) show that S-corporations exhibit relatively fast NOL deduction usage. While small size may be associated with low realized R&D tax shields, higher age, lower loss frequency and faster NOL usage may indicate higher R&D tax shields. Accordingly, our results do not cleanly map to S-corporations.

6 Conclusion

The United States' federal tax code subsidizes corporate R&D through an income tax deduction and a nonrefundable tax credit, but statutory parameters alone do not determine the economic value of these provisions because the corporate tax code is asymmetric with respect to profits and losses. Using confidential administrative tax data, we provide the first firm-level estimates of the realized present value of R&D tax benefits by tracking the generation and utilization of R&D-related deductions and credits over time, respecting rules governing carrybacks, carryforwards, and FIFO queuing across vintages.

We document substantial heterogeneity in both the timing and extent of utilization. Firms in loss status cannot fully monetize the R&D deductions and credits generated by their annual R&D spending and instead rely on delayed carryforward use. As a result, realized present-value R&D tax shields are systematically lower for young, small firms that are more likely to experience frequent losses. For the average firm in our panel, $1 of R&D investment generates about $0.36 in realized present-value tax benefits, but the realized benefit per dollar falls sharply for loss firms and is especially low for young, small firms in loss status. Thus, the firms whose expenses tilt most heavily toward R&D are often those least able to convert statutory incentives into realized tax benefits.

Our findings have two main implications. First, they speak directly to the measurement of corporate R&D tax incentives in empirical work. Because present-value tax shields vary systematically with firm state — loss status, age, and size — user-cost measures constructed solely from statutory parameters can mismeasure the effective incentives firms face. At the same time, the present-value tax shields we estimate are ex post outcomes of endogenous taxable-income and utilization paths rather than forward-looking incentive measures that govern investment decisions. The realized net acquisition price wedges we estimate therefore discipline, but do not replace, the relevant ex ante incentive objects. An open question is how to construct forward-looking asymmetry-adjusted user-cost measures that incorporate expected realization wedges conditional on firm state, particularly when investment activity simultaneously generates both a deduction and a credit.

Second, our results inform the design of reforms such as credit refundability. Making credits refundable accelerates tax benefits that would otherwise be realized only through delayed carryforward usage for firms with little or no current tax liability, including firms in loss status and firms with sufficient NOL carryforwards to eliminate current taxable income, thereby increasing their realized R&D tax shield. However, because large and incumbent firms account for the majority of R&D credit generation, much of the static dollar cost of broad refundability accrues to firms that already generate and utilize most of the credit relatively quickly. This tension highlights a central design tradeoff: while refundability can alleviate the constraints that limit small firms' usage of their R&D tax benefits, it does so at substantial cost as large fiscal benefits accrue to firms that face relatively small realization wedges. Alternative policies in the United Kingdom (Dechezleprêtre et al., 2023) and Canada (Agrawal et al., 2020) limit refundability to firms below size thresholds. Refundable designs may also create compliance challenges: the government of the United Kingdom has warned that the R&D credit for Small and Medium Enterprises ("SME") is a target for abuse and fraudulent claiming (HM Revenue & Customs, 2021, 2023).

More broadly, our evidence shows that incomplete utilization is not an innocuous accounting detail but a first-order feature of R&D tax policy in the United States. The gap between statutory and realized tax benefits shapes the distribution of incentives across firms. Any evaluation of R&D tax policy — whether focused on investment responsiveness, distributional consequences, or revenue scoring — must contend with tax asymmetries and the dynamic realization of tax attributes.

References

Aghion, P. and Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, 60:323–351.

Agrawal, A., Rosell, C., and Simcoe, T. (2020). Tax credits and small firm r&d spending. *American Economic Journal: Economic Policy*.

Altshuler, R. (1988). A dynamic analysis of the research and experimentation credit. *National Tax Journal*.

Altshuler, R. and Auerbach, A. J. (1990). The significance of tax law asymmetries: An empirical investigation. *Quarterly Journal of Economics*.

Auerbach, A. J. (1983). Corporate taxation in the united states. *Brookings Papers on Economic Activity*, 14(2):451–514.

Auerbach, A. J. (1986). The dynamic effects of tax law asymmetries. *The Review of Economic Studies*, 53(2):205–225.

Auerbach, A. J. and Poterba, J. M. (1987). Tax loss carryforwards and corporate tax incentives. In Feldstein, M., editor, *The Effects of Taxation on Capital Accumulation*, chapter 10, pages 305–342. University of Chicago Press, Chicago.

Bloom, N., Schankerman, M., and Van Reenen, J. (2013). Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4):1347–1393.

Bloom, N., Van Reenen, J., and Williams, H. (2019). A toolkit of policies to promote innovation. *Journal of Economic Perspectives*, 33(3):163–184.

Chodorow-Reich, G. (2025). The neoclassical theory of firm investment and taxes: A reassessment. Working Paper 33922, National Bureau of Economic Research.

Chodorow-Reich, G., Smith, M., Zidar, O. M., and Zwick, E. (2024). Tax policy and investment in a global economy. Working Paper 32180, National Bureau of Economic Research.

Dechezleprêtre, A., Einiö, E., Martin, R., Nguyen, K.-T., and Van Reenen, J. (2023). Do tax incentives increase firm innovation? an RD design for R&D, patents, and spillovers. *American Economic Journal: Economic Policy*, 15(4):486–521.

Eisfeldt, A. L., Hartman-Glaser, B., Kim, E. T., and Lee, K. B. (2026). Intangible intensity. Working Paper 34882, National Bureau of Economic Research.

Goodman, L., Patel, E., and Saunders-Scott, M. (2023). Implications of tax loss asymmetry for owners of S corporations. *American Economic Journal: Economic Policy*, 15(1):342–369.

Hall, R. E. and Jorgenson, D. W. (1967). Tax policy and investment behavior. *American Economic Review*.

Hatch, O. G. (2002). Statement on the research and experimentation tax credit. *Congressional Record*, 148:S170–S171. Daily ed.; statement of Sen. Hatch.

HM Revenue & Customs (2021). Preventing abuse of research and development tax relief for small and medium-sized enterprises. Policy paper, HM Revenue & Customs.

HM Revenue & Customs (2023). Hmrc annual report and accounts 2022 to 2023. Corporfate report, HM Revenue & Customs. Performance analysis section.

Joint Committee on Taxation (2011). Tax incentives for research, experimentation, and innovation. Technical Report JCX-45-11, Joint Committee on Taxation.

Joint Committee on Taxation (2015). Technical explanation of the revenue provisions of the protecting americans from tax hikes act of 2015, house amendment #2 to the senate amendment to h.r. 2029 (rules committee print 114-40). Technical Report JCX-144-15, Joint Committee on Taxation.

Joint Committee on Taxation (2018). General explanation of public law 115-97. Technical Report JCS-1-18, Joint Committee on Taxation.

Joint Committee on Taxation (2025a). Overview of the federal tax system as in effect for 2025. Technical Report JCX-38-25, Joint Committee on Taxation.

Joint Committee on Taxation (2025b). Revenue estimating process. Technical report, Joint Committee on Taxation.

Rao, N. (2016). Do tax credits stimulate r&d spending? the effect of the r&d tax credit in its first decade. *Journal of Public Economics*, 140:1–12.

Rao, N. L. and Simcoe, T. (2025). Tax incentives for research and development: Policy design and evidence. In Jones, B. and Lerner, J., editors, *Entrepreneurship and Innovation Policy and the Economy, Volume 5*, chapter 3. University of Chicago Press. NBER Chapters.

Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy*, 98(5, Part 2):S71–S102.

Zwick, E. (2021). The costs of corporate tax complexity. *American Economic Journal: Economic Policy*, 13(2):467–500.

Online Technical Appendix:

Estimating the Present Value of R&D Tax Benefits in the United States

Brandon Pecoraro, Nicholas Hoffman, Martin Lopez-Daneri,
Elena Derby, Rachel Moore, Shannon Sledz

A Estimating Observed Usage Patterns

A.1 Identifying R&D Expenditures From Tax Data

During our sample period 2010–2017, §174 of the Internal Revenue Code (the "Code") provided that research and development ("R&D") expenditures may be deducted in the year paid or incurred.[1] Qualifying R&D expenditures generally include costs incurred for activities intended to discover information that would eliminate uncertainty concerning the development or improvement of a product. Unlike deductions for wage expenses, for example, the R&D deduction is reported as part of "other deductions" on Form 1120, and therefore cannot be identified in isolation.

The Code also provides a credit for certain R&D expenses under §41. Relative to the deduction, the set of expenses that qualify for the credit is relatively narrow and composed of three groups: (i) Certain energy consortia payments, (ii) basic research payments in excess of a fixed base amount, and (iii) qualified research expenditures ("QRE") in excess of a fixed base amount. The "Regular Credit" provides for a 20% credit for the amounts in each category, where the base amount for QRE is a fixed portion of a firm's average annual gross receipts, or 50% of their QRE, whichever is smaller. Firms may elect an "Alternative Simplified Credit" that reduces the credit rate to 14% of their QRE in excess of 50% of their average QRE over the prior three years.[2] Qualifying expenses are reported on Form 6765.

The set of R&D expenses that qualify for the credit also generally qualify for the deduction, subject to the §280c election that limits the double benefit.[3] Since the R&D expenses that generate the credit and identify a subset of expenses that generate the deduction, we limit our focus to R&D expenditures reported on Form 6765 by summing basic research payments, energy consortia payments, and QRE. Since an additional set of observable R&D expenses may also qualify for the deduction for any given firm, our estimates for the amount of R&D deductions generated in any year should be interpreted as a lower bound.[4]

[1] Taxpayers may elect instead to capitalize and amortize such domestic expenditures over a period of not less than 60 months. For tax years beginning after 2021 and before 2025, capitalization of specified R&D expenditures was mandatory under modifications made in P.L. 115-97, with amortization over five years for domestic research and fifteen years for foreign research (Joint Committee on Taxation, 2018). For tax years beginning after 2024, P.L. 119-21 restored current deductibility of domestic R&E expenditures through new §174A. *Foreign* R&D expenditures, which are attributable to research conducted outside the US, must currently be capitalized and amortized over 15 years under §174. See also Joint Committee on Taxation (2025a) for a present-law summary.

[2] Qualified startups may also use the credit to offset up to $500,000 in payroll taxes. This amount was increased from $250,000 by enactment of P.L. 117-169.

[3] For a corporation that claims both the §41 credit and the §174 deduction, §280c limits the maximum combined tax benefit. Specifically, taxpayers must reduce the deduction by the amount of R&D expenses taken into account for the credit, or elect to reduce credit by the tax value of R&D expenditures deducted.

[4] While we do not directly observe deductions for R&D expenses on Form 1120, we do observe elections to amortize deductible R&D expenses on Form 4562. We drop these firms from our sample because they may confound our deduction imputations. These firms represent less than 0.2% of sample-weighted firm-cohort observations.

A.2 §174 R&D Deduction

A.2.1 Year-1 Identities: "New-Vintage" Cohort and "Old-Vintage" Stock of NOLs

Table A1 lists variables observed directly on tax returns that are relevant for estimating deduction utilization, with 2017 as the reference year for Forms 1120 and 6765. In what follows, we omit firm subscripts i to reduce notational clutter.

Table A1: Observable Tax Form Variables: Form 1120 and Form 6765 (2017)

Variable	*Description*	*Source*
$\texttt{exp}_t^{RD}$	R&D Expenses	Form 6765, Section A (B), lines 1+9 (18+28)
$\texttt{crd}_t^{RD}$	R&D credit	Form 6765, Section A (B) line 17 (34) box
ρ_t	280c election	Form 6765, Section A (B), line 17 (34)
$\texttt{inc}_t$	Total income	Form 1120, line 11
$\texttt{ded}_t$	Total deductions (excluding NOL and special deductions)	Form 1120, line 27
$\texttt{TI}_t^b$	Taxable income *before* NOL and special deductions	Form 1120, line 28
$\texttt{NOL}_t$	Preexisting NOL carryforwards at start of year t	Form 1120, Schedule K, line 11
$\texttt{use}_t^{NOL}$	NOL deductions	Form 1120, line 29a
$\texttt{TI}_t$	Taxable income	Form 1120, line 30

Let TI_t^b be taxable income computed before NOL and special deductions. If $\text{TI}_t^b > 0$, then the corporation is in non-loss status and fully utilizes all deductions claimed as part of its total deductions in the current year (excluding NOL and special deductions), which includes the R&D deduction. If instead $\text{TI}_t^b < 0$, then the corporation is in loss status and may not fully utilize all deductions claimed in the current year. For a loss-status corporation, the amount of total deductions in excess of total income is defined as the loss, and may be carried forward or carried back (when the Code permits). Generally:

$$loss_t \equiv \max(0, -\texttt{TI}_t^b) = CB_t^{NOL} + CF_t^{NOL} \tag{1}$$

where CB_t^{NOL} is the amount (if any) of the current-year loss carried back as a NOL deduction against a prior year's income and CF_t^{NOL} is the amount (if any) of the current-year loss carried forward to potentially be used as a NOL deduction against future income.

At the beginning of each year, a corporation may have a preexisting stock of NOL carryforwards generated by past losses. This observed "old-vintage" stock, $\texttt{NOL}_1$, will be reduced when the corporation claims a current-year NOL deduction or negative "outside" adjustments occur such as expiration.[5] "New-vintage" carryforwards of current-year losses will increase the stock of NOL carryforwards brought into the subsequent year. The law of motion for the stock of a corporation's total NOL carryforward stock across any two years is:

$$\texttt{NOL}_{t+1} = \texttt{NOL}_t - \texttt{use}_t^{\text{NOL}} + CF_t^{NOL} + adj_t^{\text{NOL}} \tag{2}$$

where $\texttt{use}_t^{\text{NOL}}$ is the amount of the NOL stock observed to be used to reduce current taxable income, adj_t^{NOL} is negative or positive change to the subsequent year's stock of NOLs that may arise due to expiration of NOLs or other "outside" adjustments.

For any initial year $t = 1$, equations (1) and (2) represent three unknowns and two equations. In order to jointly impute the unknowns $\{CB_1^{NOL}, CF_1^{NOL}, adj_1^{\text{NOL}}\}$, assumptions may be used along with the accounting identities according to the following procedure:

1. If election is made to forgo the carryback period, then current-year carryforwards are the full amount of the current loss:

$$\begin{aligned} CB_1^{NOL} &= 0 \\ CF_1^{NOL} &= loss_1 \end{aligned} \tag{3}$$

   If the election to forgo the carryback period is not made, then current-year losses are first used to reduce taxable income (if any) in the previous two tax years, and any remaining losses are carried forward:

$$\begin{aligned} CB_1^{NOL} &= \min\{loss_1, \sum_{i=1}^{2} \texttt{TI}_{-i}\} \\ CF_1^{NOL} &= loss_1 - CB_1^{NOL} \end{aligned} \tag{4}$$

2. Define the portion of the change in the stock of NOLs that can be observed in the data as:

$$dNOL_1 \equiv \texttt{NOL}_2 - (\texttt{NOL}_1 - \texttt{use}_1^{\text{NOL}}) = CF_1^{NOL} + adj_1^{\text{NOL}} \tag{5}$$

   which leaves adj_1^{NOL} identified as the residual.

3. If the corporation has an old-vintage stock of NOL carryforwards so that $\texttt{NOL}_1 > 0$, it must be exhausted before any of the new-vintage carryforwards associated with the year-1 R&D deduction can be used.[6] After any NOL usage, the remaining amount of old-vintage NOL

[5] "Outside" adjustments may also arise from mergers and acquisitions subject to §381, changes to consolidation groups on consolidated tax returns, prior-year amended returns, IRS exam adjustments, changes to prior-year carrybacks, or accounting-method changes.

[6] The FIFO convention renders the old-vintage and new-vintage NOL stocks as non-fungible since such carryforwards expire after 20 years. In the absence of carryforward expiration, changing tax rates, or other nonlinear constraints, FIFO affects the measured realization of particular vintages but not the underlying marginal present value of tax attributes.

stock is defined as:

$$rem_2^{oldNOL} \equiv \max\{0,\ \texttt{NOL}_1 - \texttt{use}_1^{NOL} + adj_1^{\text{NOL}}\} \tag{6}$$

Outside adjustments that arise from reasons unrelated to carryforward expiration or prior-year amended returns may introduce noise to our estimates. For purposes of tracking old-vintage NOL usage, we clip outliers on the adj_t^{NOL} term that exceed 1.5 interquartile ranges above the upper quartile or below the lower quartile.

Isolating the R&D Deduction: Since the §174 R&D deduction is contained in the total deductions used, carried back, or carried forward, it may be imputed once each unknown is identified. Let $\texttt{exp}_1^{RD}$ and $\texttt{crd}_1^{RD}$ be the total amount of R&D expenses and the total credit claimed, respectively, as observed on Form 6765 for a given year one. Letting ρ_1 be equal to 1 if the firm makes the §280c election to take the full credit in year one, we impute R&D deductions as:[7]

$$ded_1^{RD} = \texttt{exp}_1^{RD} - \rho_1 \texttt{crd}_1^{RD}$$

Since no particular stacking order is specified in the Code for each of the individual deductions, the amount of usage, carryback, or carryforward attributed to any individual deduction is assumed to be proportional to its share of total deductions, $\texttt{ded}_1$. Therefore, first-year usage of the R&D deduction can be expressed as the amount claimed less the attributable portion of new losses carried forward in excess of carrybacks (if any):

$$use_1^{\text{RD},ded} \equiv \begin{cases} ded_1^{RD} & \text{if } \texttt{TI}_1^{\text{b}} \geq 0 \\ ded_1^{RD} - \underbrace{(loss_1 - CB_1^{NOL})}_{CF_1^{NOL}} \left(\frac{ded_1^{RD}}{\texttt{ded}_1}\right) & \text{if } \texttt{TI}_1^{\text{b}} < 0 \end{cases} \tag{7}$$

A.2.2 Years 2 through N: Accounting for Observed Usage

A NOL deduction may be used to reduce year-t taxable income when $\text{TI}_t^b > 0$ and $\texttt{NOL}_t > 0$. Under the FIFO timing convention, old-vintage NOL stock must be completely exhausted before any new-vintage NOL carryforwards generated in year one may be used. The procedure below is used to account for usage of these two different vintages of NOL carryforward stocks while isolating the portion of the new vintage that is attributable to the R&D deduction.

For $t \in \{2, \ldots, N\}$:

1. If there remains any old-vintage NOL stock at the beginning of year t so that $rem_t^{oldNOL} > 0$, observed NOL deductions are attributed to the old vintage so that:

[7] Since we identify R&D expenses as those that generate the credit, a firm that elects to take the full R&D credit will be assigned zero R&D deductions. In our sample, only about 25% of sample-weighted firm-cohort observations elect to take the full credit.

$$use_t^{oldNOL} \equiv \min\left(\texttt{use}_t^{NOL}, rem_t^{oldNOL}\right) \tag{8}$$

Since the remaining stock of the old-vintage NOL carryforwards may also be reduced via expiration or increased via other negative "outside" adjustments, the law of motion for the old-vintage NOL carryforward stock can be expressed as:

$$rem_{t+1}^{oldNOL} = \max\{0, rem_t^{oldNOL} - use_t^{oldNOL} + adj_t^{\text{NOL}}\} \tag{9}$$

2. Define the remaining amount of new-vintage NOLs generated in year one attributable to the R&D deduction as:

$$rem_t^{RD,ded} \equiv ded_1^{RD} - \sum_{s<t} use_s^{RD,ded} \tag{10}$$

When the observed year-t NOL deduction exceeds the remaining old-vintage NOL stock, any positive difference is attributed to usage of the new-vintage NOL carryforwards. If the old vintage is fully exhausted, then $use_t^{oldNOL} = 0$ and the full amount of any observed NOL deductions are attributed to usage of new-vintage NOL carryforwards. Any usage of the new vintage is distributed proportionally to all year-1 individual deductions so that the amount attributable to the R&D deduction can be expressed as:

$$use_t^{RD,ded} = \min\left\{(\texttt{use}_t^{NOL} - use_t^{oldNOL})\left(\frac{ded_1^{RD}}{\texttt{ded}_1}\right), \quad rem_t^{RD,ded}\right\} \tag{11}$$

For a firm with $rem_{N+1}^{oldNOL} = 0$, any old-vintage NOL carryforwards have been fully exhausted within the sample period, and usage has begun on any new-vintage NOL carryforwards generated by the year-1 deduction. For a firm with $rem_{N+1}^{RD,ded} = 0$, any portion of new-vintage NOL carryforwards attributable to the year-1 R&D deduction are also fully exhausted within the sample period. In such cases, the observed sequence $\{use_1^{RD,ded}, ..., use_N^{RD,ded}\}$ fully identifies the present-value R&D tax shield.

A.3 §41 RD Credit

A.3.1 Year-1 Identities: "New-Vintage" Cohort and "Old-Vintage" Stock of GBCs

Table A2 lists variables observed directly on tax returns that are relevant for estimating credit utilization, with 2017 as the reference year for Forms 6765 and 3800. In what follows we restrict our attention to "non-specified credits," which is the subset of current-year credits that contains the §41 R&D credit, and are not allowable against Alternative Minimum Tax ("AMT") liability.[8] We omit firm subscripts i to reduce notational clutter.

[8] Non-specified credits contrast with "specified" credits, which are allowable against AMT liability. While the R&D credit qualifies as specified credit for "eligible small businesses" whose stock is not publicly traded, our sample of corporations does not include any ESB.

Table A2: Additional Variables Used for Credits

Variable	*Description*	*Source*
$\texttt{crd}_t^{RD}$	RD credit generated in year one	Form 6765
$\texttt{crd}_t$	New-vintage (non-specified) GBCs not allowable against TMT	Form 3800, Part I, lines 1 + 3
$\texttt{GBC}_t$	Preexisting stock of GBC carryforwards not allowed against TMT	Form 3800, Part I, line 5
$\texttt{T}_t^{inc}$	Net Income Tax *before* GBC	Form 3800, Part II, line 11
$\texttt{T}_t^{reg}$	Net Regular Tax *before* GBC	Form 3800, Part II, line 12
$\texttt{TMT}_t$	Tentative Alternative Minimum Tax	Form 3800, Part II, line 14
$\texttt{L}_t^{GBC}$	GBC limitation	Form 3800, Part II, line 16
$\texttt{use}_t^{totGBC}$	Total (specified and non-specified) GBCs used	Form 3800, Part II, line 38

The total amount of non-specified GBCs used in any year t is limited to:

$$\texttt{L}_t^{GBC} \equiv \texttt{T}_t^{inc} - \max\left\{\texttt{TMT}_t,\ 0.25 \times (\texttt{T}_t^{reg} - \$25,000)\right\} \tag{12}$$

which is the excess of a corporation's net income tax over the greater of the tentative minimum tax or 25%. Computed without regard to carrybacks, the total amount of GBC usage in any year respects the following identity:

$$use_t^{GBC} = \min\{\texttt{GBC}_t + \texttt{crd}_t, \texttt{L}_t^{GBC}\} \tag{13}$$

The variable use_t^{GBC} may contain both old-vintage GBC carryforwards and current-year credits, the former of which must be used first under the FIFO rules of §38 of the Code. Exploiting the FIFO convention, usage of each can be imputed as follows:

$$use_t^{oldGBC} = \min\{\texttt{GBC}_t,\ \texttt{L}_t^{GBC}\} \tag{14}$$

$$use_t^{crd} = use_t^{GBC} - use_t^{oldGBC} \tag{15}$$

If the GBC limitation is not binding such that $\texttt{GBC}_t + \texttt{crd}_t < \texttt{L}_t^{GBC}$, then the corporation fully utilizes all old-vintage GBC carryforwards as well as all current-year credits so that $use_t^{oldGBC} = \texttt{GBC}_t$ and $use_t^{crd} = \texttt{crd}_t$. If instead the GBC limitation is binding, then the corporation will not fully utilize all current-year credits. The portion of current-year credits not used to reduce tax liability becomes a GBC carryover. This carryover must first be used as a carryback against the prior year's tax liability, CB_t^{GBC}, with any residual becoming a "new-vintage" carryforward to be used against a future year's tax liability, CF_t^{GBC}. We thus have the following identity for current-year credits:

$$\texttt{crd}_t \equiv use_t^{crd} + CB_t^{GBC} + CF_t^{GBC} \tag{16}$$

Given the statutory ordering rule for each individual current-year credit as specified in §38 of the Code, usage of any current-year R&D credit $\texttt{crd}_t^{RD} \in \texttt{crd}_t$ can be imputed from overall GBC usage.

Since old-vintage GBC carryforwards are applied before current-year credits, the observed stock of old-vintage GBC carryforwards (if any) will be reduced when the corporation claims a current-year GBC. This stock of GBC carryforwards can also decrease due to "outside" adjustments, adj_t^{GBC}.[9] Incomplete usage of current-year credits will create new-vintage GBC carryforwards that increase the corporation's total stock of GBC carryforwards, as will positive outside adjustments. The law of motion for the stock of a corporation's GBC carryforward stock across any two years is:

$$\texttt{GBC}_{t+1} = \texttt{GBC}_t - use_t^{oldGBC} + CF_t^{GBC} + adj_t^{GBC}. \tag{17}$$

For the initial year $t = 1$, usage variables $\{use_1^{oldGBC}, use_1^{crd}\}$ are jointly determined by equations (14) and (15). Equations (16) and (17) represent three unknowns in two equations. In order to jointly impute the unknowns $\{CB_1^{GBC}, CF_1^{GBC}, adj_1^{GBC}\}$, assumptions may be used along with the accounting identities according to the following procedure:

1. The unused current-year credits are first carried back, with any remaining portion generating a new-vintage GBC carryforward:

$$\begin{aligned} CB_1^{GBC} &= \min\Big\{\texttt{crd}_1 - use_1^{crd},\, \texttt{L}_{-1}^{GBC} - use_{-1}^{GBC}\Big\} \\ CF_1^{GBC} &= \texttt{crd}_1 - use_1^{crd} - CB_1^{GBC} \end{aligned} \tag{18}$$

2. Define the change in the stock of GBCs that can be observed in the data as:

$$dGBC_1 \equiv \texttt{GBC}_2 - \big(\texttt{GBC}_1 - use_1^{oldGBC}\big) = CF_1^{GBC} + adj_1^{GBC}. \tag{19}$$

   which leaves adj_1^{GBC} identified as the residual.

3. If the corporation has any old-vintage stock of GBC carryforwards so that $\texttt{GBC}_1 > 0$, it must be exhausted before any of the new-vintage GBC carryforwards can be used. Accounting for carryforward usage, the remaining amount of the old-vintage GBC stock is defined as:

$$rem_2^{oldGBC} \equiv \max\{0,\ \texttt{GBC}_1 - \texttt{use}_1^{oldGBC} + adj_1^{GBC}\} \tag{20}$$

Outside adjustments that arise from reasons unrelated to carryforward expiration or prior-year amended returns may introduce noise to our estimates. For purposes of tracking old-vintage GBC usage, we clip outliers on the adj_t^{GBC} term that exceed 1.5 interquartile ranges above the upper quartile or below the lower quartile.

Isolating the R&D Credit: The R&D credit can be isolated from the other current-year credits in the GBC regime by partitioning $\texttt{crd}_1$ into three pieces based on the within-period credit ordering

[9] As with the case of the stock of NOL carryforwards, such "outside" adjustments can occur from expiration of GBC carryforwards, exam adjustments, mergers, or released credits.

rule specified in §38 of the Code:

$$\texttt{crd}_1 = \underbrace{\texttt{crd}_1^{pre}}_{\text{sum of credits ordered before RD}} + \underbrace{\texttt{crd}_1^{RD}}_{\text{§41 RD credit}} + \underbrace{\texttt{crd}_1^{post}}_{\text{sum of credits ordered after RD}}.$$

For purposes of measuring utilization of the R&D credit, the credits contained in $\texttt{crd}_1^{pre}$ must be completely used against net income tax liability before the R&D credit can be considered used. Under the credit ordering rule, usage of these preceding credits can be defined without regard to carrybacks as:

$$use_1^{pre} \equiv \min\{\texttt{crd}_1^{pre},\ use_1^{crd}\}. \tag{21}$$

The positive difference $use_1^{crd} - use_1^{pre}$ (if any) then defines the residual capacity for the R&D credit to be used in the first year:

$$use_1^{RD,crd} \equiv \min\{\texttt{crd}_1^{RD},\ use_1^{crd} - use_1^{pre}\}, \tag{22}$$

where the amount of $\texttt{crd}_1^{post}$ used can be obtained as a residual.

Next, the carryback amount (if any) must be allocated:

$$CB_1^{pre} \equiv \min\{\texttt{crd}_1^{pre} - use_1^{pre},\ CB_1^{GBC}\} \tag{23}$$

$$CB_1^{RD} \equiv \min\{\texttt{crd}_1^{RD} - use_1^{RD,crd}, \max\{CB_1^{GBC} - CB_1^{pre}, 0\}\}, \tag{24}$$

With slight abuse of notation, the amount of current-year R&D credit usage must be updated to account for any carrybacks:

$$use_1^{RD,crd} \leftarrow use_1^{RD,crd} + CB_1^{RD}$$

A.3.2 Years 2 through N: Accounting for Observed Usage

Under the FIFO convention, any old-vintage stock of GBC carryforwards must be exhausted before any of the new-vintage GBC carryforwards may be used. The procedure below is used to account for usage of these two different vintages of GBC carryforwards while isolating the portion of the new vintage that is attributable to the R&D credit.

For $t \in \{2, \ldots, N\}$:

1. If there remains any old-vintage GBC stock so that $rem_t^{oldGBC} > 0$, observed usage of GBC carryforwards is attributed to the old vintage so that:

$$use_t^{oldGBC} \equiv \min\{rem_t^{oldGBC},\ use_t^{GBC}\} \tag{25}$$

with the law of motion for old-vintage GBC carryforwards defined in a way analogous to the NOL case:

$$rem_{t+1}^{oldGBC} = \max\Big\{0,\ rem_t^{oldGBC} - use_t^{oldGBC} + \min\{0, adj_t^{GBC}\}\Big\}, \tag{26}$$

2. We separately track the remaining new-vintage carryforwards associated with the year-1 credit that precedes the R&D credit in the ordering rule, and the remaining amount of the new-vintage carryforward associated with R&D credit itself.

$$rem_t^{pre} \equiv \texttt{crd}_1^{pre} - \sum_{s<t} use_s^{pre} \tag{27}$$

$$rem_t^{RD,crd} \equiv \texttt{crd}_1^{RD} - \sum_{s<t} use_s^{RD,crd} \tag{28}$$

When the observed year-t total GBC carryforward usage exceeds the remaining old-vintage GBC carryforward stock, the excess is attributed to usage of the new-vintage GBC carryforwards. If the old vintage is fully exhausted, then $use_t^{oldGBC} = 0$ and the full amount of any observed GBC carryforward usage is attributed to the new vintage. Any usage of the new vintage follows the statutory credit ordering rule so that the new-vintage carryforwards attributed to the credit taking priority over the R&D credit are used first:

$$use_t^{pre} \equiv \min\{rem_t^{pre},\ use_t^{GBC} - use_t^{oldGBC}\}. \tag{29}$$

The positive difference $use_t^{GBC} - use_t^{oldGBC} - use_t^{pre}$ (if any) defines the residual capacity for usage of new-vintage carryforward associated with the R&D credit:

$$use_t^{RD,crd} \equiv \min\{rem_t^{RD,crd},\ use_t^{GBC} - use_t^{oldGBC} - use_t^{pre}\} \tag{30}$$

For a firm with $rem_{N+1}^{oldGBC} = 0$, any old-vintage GBC carryforwards have been fully exhausted within the sample period, and usage has begun on any new-vintage GBC carryforwards. For a firm with $rem_{N+1}^{RD,crd} = 0$, any portion of the new-vintage GBC carryforward attributable to the R&D credit is also fully exhausted within the sample period. In such cases, the observed sequence $\{use_1^{RD,crd}, \ldots, use_N^{RD,crd}\}$ fully identifies the credit portion of the present-value R&D tax shield.

B Extrapolating Unobserved Usage Patterns

B.1 Estimating Burn Rates

For firms with remaining new-vintage NOL and/or GBC carryforwards in year $N+1$, the unobserved portion of usage needs to be extrapolated in order to compute the present-value objects. Our procedure for doing so involves the estimation of a bin-specific "burn rate" for the remaining stock of old-vintage and new-vintage carryforwards, which will be used for forward extrapolation. Let

j index the number of bins $J = 4$. For purposes of extrapolating old-vintage NOL carryforwards, we bin firms by the number of years from $t = 1$ through $t = N$ in which they are in loss status, separating firms with zero, one, two, and three-or-more loss years. For purposes of extrapolating old-vintage GBC carryforwards, we bin firms by the number of years from $t = 1$ through $t = N$ in which they are GBC-limited, separating firms with zero, one, two, and three-or-more constrained years. For purposes of new-vintage carryforward usage, we do not bin firms.

We estimate burn rates by pooling all observed years 2012-2016. In what follows the superscript "old" denotes the relevant old-vintage attribute while the superscript "RD" denotes the relevant new-vintage attribute (either R&D deduction or credit component), and we apply the procedure separately to deductions and credits:

1. **Burn Rates for Old-Vintage Carryforwards**: This governs the *queue* until the new-vintage carryforwards can be used. For each j bin:

$$\widehat{b}(j)^{\text{old}} = \frac{\sum\limits_{i,t:\, rem(j)_{i,t}^{old}>0} \min\{use(j)_{i,t}^{old},\, rem(j)_{i,t}^{old}\}}{\sum\limits_{i,t:\, rem(j)_{i,t}^{old}>0} rem(j)_{i,t}^{old}} \tag{31}$$

   Our estimates are $\widehat{b}(j)^{\text{oldNOL}} = \{0.36, 0.18, 0.16, 0.05\}$ and $\widehat{b}(j)^{\text{oldGBC}} = \{0.15, 0.12, 0.15, 0.02\}$. We also estimate a pooled-sample burn rate without bins, $\{\widehat{b}_{\text{all}}^{\text{oldNOL}}, \widehat{b}_{\text{all}}^{\text{oldGBC}}\} = \{0.14, 0.05\}$.

2. **Burn Rate for New-Vintage carryforwards:** This governs *how fast* the new-vintage carryforwards are used once the old-vintage carryforwards are exhausted. Letting $x_{i,1}^{RD}$ stand in for $ded_{i,1}^{RD}$ and $crd_{i,1}^{RD}$ respectively:

$$\widehat{b}^{\text{RD}} = \frac{\sum\limits_{i,t:\, rem_{i,t}^{old}=0} use_{i,t}^{RD}}{\sum\limits_{i,t:\, rem_{i,t}^{old}=0} \left(x_{i,1}^{RD} - \sum\limits_{s<t} use_{i,s}^{RD}\right)} \tag{32}$$

   We estimate $\{\widehat{b}_{\text{all}}^{\text{dedRD}}, \widehat{b}_{\text{all}}^{\text{crdRD}}\} = \{0.24, 0.17\}$.

B.2 Forward Simulation of Usage Using Burn Rates

Given $\widehat{b}(j)^{\text{old}}$, $\widehat{b}_{\text{all}}^{\text{old}}$, and $\widehat{b}^{\text{RD}}$, we extrapolate the unobserved usage path forward from $t = N+1, \ldots, 21$. The simulation treats the remaining stocks at $t = N + 1$ as initial conditions and iterates forward deterministically, applying the estimated burn rates to the remaining stock each period. Throughout, old-vintage carryforwards are used first; new-vintage carryforwards can only be used once the old-vintage queue is deemed exhausted under FIFO rules. To avoid an asymptotic tail for old-vintage carryforwards, we impose for each firm i an old-vintage exhaustion threshold $f_i^{\text{old}} \equiv 0.01\, rem_{i,N+1}^{\text{old}}$, and treat the stock as exhausted whenever it falls below its corresponding floor. Once the simulation reaches $t = 20$, any remaining old-vintage carryforwards must expire.[10] The simulation is terminated at $t = 21$ when any remaining new-vintage carryforwards must expire.

While each firm is assigned its bin-specific burn rate in the first year of the extrapolation procedure, we allow for gradual convergence in burn rates across firms via mean-reversion. Let $b_i^{\text{old}} \equiv \widehat{b}(j)^{\text{old}}$ for $i \in j$ denote the old-vintage burn rate assigned to a given firm via the binning process and $\widehat{b}_{i,t}^{\text{old}}$ denote the burn rate used by the firm in year $t \geq N+1$ of the simulation. We then specify the mean-reversion process as:

$$\widehat{b}_{i,t}^{\text{old}} \;=\; \widehat{b}_{\text{all}}^{\text{old}} \;+\; \lambda^{t-(N+1)}\big(b_i^{\text{old}} - \widehat{b}_{\text{all}}^{\text{old}}\big),$$

so that cross-sectional heterogeneity in burn behavior decays geometrically at rate $\lambda \in (0,1)$. We estimate λ using a de-meaned AR(1) specification on firm-level burn rates over the pooled series 2012-2016, including zero-use years. In our simulations we use $\lambda^{\text{NOL}} = 0.17$ and $\lambda^{\text{GBC}} = 0.47$.

For each $t \geq N+1$, the simulation proceeds in two steps (applied separately to NOL and GBC carryforwards):

1. **Old-Vintage Usage:** If $rem_{i,t}^{\text{old}} > f_i^{\text{old}}$, old-vintage usage is given by

   $$use_{i,t}^{\text{old}} \;=\; \widehat{b}_{i,t}^{\text{old}}\, rem_{i,t}^{\text{old}},$$

   with the implied update

   $$rem_{i,t+1}^{\text{old}} \;=\; rem_{i,t}^{\text{old}} - use_{i,t}^{\text{old}}.$$

   If instead $rem_{i,t}^{\text{old}} \leq f_i^{\text{old}}$, we set $use_{i,t}^{\text{old}} = rem_{i,t}^{\text{old}}$ and $rem_{i,t+1}^{\text{old}} = 0$.

2. **New-vintage Usage:** New-vintage carryforwards are used only after the old-vintage queue is exhausted. Accordingly, if $rem_{i,t+1}^{\text{old}} = 0$ then:

   $$use_{i,t}^{\text{RD}} \;=\; \widehat{b}^{\text{RD}}\, rem_{i,t}^{\text{RD}},$$

   with update

   $$rem_{i,t+1}^{\text{RD}} \;=\; rem_{i,t}^{\text{RD}} - use_{i,t}^{\text{RD}}.$$

Iterating this recursion forward yields the extrapolated usage path $\{use_{i,t}^{\text{old}}, use_{i,t}^{\text{RD}}\}_{t=N+1}^{\bar{T}}$, where $\bar{T}$ is chosen sufficiently large that all remaining stocks are exhausted (up to the old-vintage floor), or expire.

Burn Rates as a Reduced-Form Extensive-Intensive Utilization Process: As described in this paper, firms' usage of carryforwards can be a discrete process with years of zero utilization

[10] Since we do not observe the age of carryforwards within the old-vintage stock, and do not perfectly observe expiration, this ensures that no old-vintage carryforwards are usable past the longest plausible horizon.

(e.g., in loss or credit-constrained years) and positive utilization in other years. While our burn-rate forward simulation procedure is specified so that firms fractionally draw down carryforward stocks each year, it can be interpreted as a reduced-form version of a two–part, extensive–intensive utilization model. In such a model, a firm draws an extensive-margin "activation" indicator that determines whether it is able to use any carryforward at all (e.g., because it is in non-loss position and not credit constrained), which occurs with probability p_t. Conditional on being active, the firm draws down a fraction $u_t \in [0, 1]$ of the remaining stock. This structure implies an expected utilization rate

$$\lambda_t \equiv p_t u_t,$$

and therefore an expected law of motion for the remaining carryforward inventory S_t,

$$S_{t+1} = (1 - \lambda_t)\, S_t.$$

Our burn-rate extrapolation adopts λ_t — the expected fraction of the remaining stock monetized each year — as the primitive object to be estimated and simulated forward. In this sense, burn rates provide a reduced-form summary of the underlying two-margin utilization process that is sufficient for our forward-simulation extrapolation of carryforward stocks.