Insurance through the Tax Code: Corporate Loss Refunds, Firm Survival, and Employment*

Ian Sapollnik
MIT

Dustin Swonder
UC Berkeley

Job Market Paper

August 26, 2026

PRELIMINARY AND INCOMPLETE -- PLEASE DO NOT CITE OR CIRCULATE

Abstract

Corporate income taxes are asymmetric: profits are taxed as they accrue, while losses are refunded only through deferred and uncertain deductions against future profits. We study Australia's introduction of tax loss carrybacks during the COVID-19 pandemic, which allowed corporations to reclaim taxes paid on past profits, using the universe of business tax returns linked to monthly payroll records. Eligibility conditioned on a history of prior tax payments and acted as a tag: eligible firms entered the loss year from long histories of profit and returned to profitability at roughly twice the rate of the average loss-maker. Take-up rose steeply with the size of the entitlement and, holding the entitlement fixed, was highest among financially constrained firms. Comparing corporations to the pass-through businesses the statute excluded, each assigned the counterfactual refund its own history would have commanded, we estimate that the refund raised three-year firm survival by 1.6 percentage points at the average entitlement, against a baseline of 85 percent, and generated persistent increases in employment and payroll, with no detectable change in pay per employee. The payroll response operates through the retention of incumbent workers rather than through new hiring. The refunds functioned as insurance for viable firms against a transitory shock, preserving employment relationships that would otherwise have dissolved.

JEL codes: H25, H32, G33, J63, E62
Keywords: tax loss carrybacks, asymmetric taxation, firm survival, corporate taxation, employment, COVID-19

*Sapollnik: ian_s@mit.edu; Swonder: dswonder@berkeley.edu. We thank Nathan Hendren, Jim Poterba and Antoinette Schoar for guidance throughout this project. We are additionally grateful to Alan Auerbach, Jack Buckley, Amy Finkelstein, Jonathan Gruber, Gianni La Cava, Rachel Lee, Adam Solomon, Evan Soltas, Danny Yagan and participants at the MIT Public Finance Lunch, MIT Applied Micro Lunch and Canadian Economics Association Annual Meeting for helpful comments and conversations. We thank the e61 Institute for their generosity in hosting Sapollnik while part of the data work for this project was completed. Sapollnik is supported in part by funding from the Social Sciences and Humanities Research Council of Canada (grant number 752-2021-0517), and gratefully acknowledges additional support from the Lincoln Institute for Land Policy. Swonder gratefully acknowledges support from a National Science Foundation Graduate Research Fellowship (grant number DGE 2146752). All analysis was carried out inside a secure environment operated by the Australian Bureau of Statistics; unit records never leave it, and all estimates and statistics in this paper were cleared for disclosure by the Bureau.

1 Introduction

Corporate income taxes are asymmetric. When a firm earns a dollar of profit, the government taxes it at once; when the firm runs a dollar of losses, the government does not symmetrically refund the loss. This asymmetry is a near-universal feature of corporate tax systems, and its consequences depend on which firms run losses and why.

Whether the government should relax the asymmetry when losses become widespread is not obvious. On one side, the asymmetry raises the effective tax rate on the activities that generate losses along the way, and thereby discourages investment and risk-taking [Auerbach, 1986, Auerbach and Poterba, 1987, Altshuler and Auerbach, 1990, Edgerton, 2010, Ljungqvist et al., 2017]; non-refundability is most consequential in economic crises, when losses spread even to productive firms and liquidity needs are greatest. On the other side, a refund flows to any firm with a loss, viable or not. Refunding losses may keep alive firms whose exit would be no loss [Caballero et al., 2008, Bethmann et al., 2018], slow the reallocation of workers and capital to better uses, and deliver windfalls to firms that would have recovered on their own, at a fiscal cost borne when budgets are most strained. Whether refunds preserve viable firms and the jobs they support, or pay for outcomes that would have occurred anyway, is an empirical question, and it is the subject of this paper.

We study a policy that made a corporate tax system more symmetric by reforming the tax treatment of losses. Under the prevailing regime, a loss-making firm could obtain relief only in future years, by *carrying forward* its loss and deducting it against later profits: relief that pays no interest while it waits, requires the firm to return to profit, and is lost if the firm does not survive. The reform allowed certain firms to instead receive immediate refunds by *carrying back* their losses, deducting them against past profits and recovering tax already paid. The refund converts a deferred and contingent claim into cash, and the cash arrives in the middle of a loss, when internal funds are scarcest. As we formalize in Section 2, the value of the conversion has two parts: a present value gain, because money now is worth more than a deduction later, and a certainty gain, because the refund is paid whether or not the firm ever returns to profit.

Refunds of past tax are rare as permanent features of tax systems but recurrent as crisis policy. The United States has extended its provision in each of the last three major downturns, in 2002, 2009 and 2020, at a cost on the order of a tenth of annual corporate tax receipts in each episode, and governments across the developed world introduced or expanded similar measures during the global financial crisis and the pandemic.[1]

Yet credible evidence on the causal effects of these measures is scarce, for a reason built into the policy itself. Eligibility depends on the firm's own history of profits and losses, so eligible firms differ from ineligible ones in exactly the ways that determine their future outcomes; and claiming is elective, so claimants differ from non-claimants in ways that reflect their present. Existing work has documented take-up and the correlates of refund use; separating what refunds cause from whom they select has been harder. The effect on workers is a further open question: a refund that saves an employer may spare its workers a spell of unemployment, or simply postpone their move to a

[1] In the United States, the Job Creation and Worker Assistance Act of 2002 extended the carryback window from two years to five for losses incurred in 2001 and 2002; the Worker, Homeownership, and Business Assistance Act of 2009 did the same for 2008 and 2009 losses; and the CARES Act of 2020 restored a five-year window for losses incurred in 2018 through 2020, the window having been eliminated entirely in 2017. The United Kingdom extended its one-year window to three years in both 2009 and 2021. Germany raised the cap on its one-year provision from EUR 1 million to EUR 10 million for 2020 and 2021 losses. France allowed 2020 losses to be carried back without its usual cap, and Japan temporarily widened eligibility for its refund from companies with capital below 100 million yen to those below 1 billion yen. New Zealand, Norway and Singapore, which ordinarily allow little or no refunding of past tax, introduced temporary or enhanced schemes in 2020.

better match.

Australia during the COVID-19 pandemic offers an unusually clean setting for these questions, for three reasons. First, the policy arrived where none existed. Australia has no permanent provision for refunding past tax, so its introduction in October 2020 changed the treatment of losses discretely rather than at the margin. Second, the statute drew a sharp line through the population of firms: only incorporated businesses could claim, while partnerships, trusts and sole traders, which operate alongside companies in every industry, were excluded whatever their losses. That line is the basis of our research design. Third, we observe the universe of business tax returns from 2001-02 through 2022-23, linked to payroll records that report every employer's workforce and wage bill monthly and in near real time. We see both the refunds and the outcomes they are meant to affect, for every employer in the economy.

We proceed in three steps. We first reconstruct, from the statutory formula and each firm's own filing history, the refund every firm was entitled to, whether or not it claimed; where claims are observed, the reconstruction reproduces them almost exactly. We then document who the policy reached. Entitlements are heavily concentrated, and the eligibility criterion is a powerful tag: eligible firms enter the loss year from long histories of profit and return to profitability at roughly twice the rate of the average loss-maker. Take-up is far from complete, echoing Zwick [2021]: only 29 percent of eligible firms claim, claiming rises steeply with the dollars at stake, and, holding the entitlement fixed, it is the financially fragile firms that claim. Finally, to separate what refunds cause from whom they select, we compare corporations to the pass-through businesses the statute excluded. For each pass-through we compute the refund it would have been entitled to had it been incorporated, match it to a corporation of similar industry, size and history, and ask whether outcomes diverge with the size of the entitlement where the entitlement was real, relative to where it was counterfactual. The comparison follows a tradition of using legal form to identify the effects of corporate taxation [Yagan, 2015, Kennedy et al., 2026], and its assumptions can be probed: outcomes of the two groups track each other in the months before refunds arrive, and entitlements constructed by the same formula in a year with no carryback produce nothing.

The refunds mattered for the firms that received them. A corporation receiving the average entitlement was roughly 1.6 percentage points more likely to be operating three years later, against a baseline of 85 percent. In monthly event studies around the expected arrival of the refund, payroll at treated firms rises by roughly ten cents per refund dollar per month and employment by roughly 0.8 workers at the average refund, with no movement in the sixteen months before arrival; average pay per employee does not detectably change. The payroll response operates through the retention of incumbent workers rather than through new hiring.

Section 2 develops an accounting framework in which the same refund dollar can act as a loan, a transfer, or a lifeline, and shows that which one it is depends on a single empirical object: the rate at which refunded losses would have been deducted in the counterfactual. Because eligibility tags firms with persistent histories of profit, much of the outlay is an advance against deductions the recipients would otherwise have taken, and operates as an interest-free loan rather than a grant. Read together, the estimates describe a policy that worked as insurance for viable firms against a transitory shock: it reached firms selected on histories that predict recovery, was claimed where cash was scarcest, and preserved employment relationships.

Our results contribute to three literatures. The first studies the asymmetric taxation of profits and losses. A theoretical tradition establishes that the asymmetry raises effective tax rates and distorts investment and risk-taking [Auerbach, 1986, Auerbach and Poterba, 1987, Altshuler and Auerbach, 1990], and empirical work documents the scale of unused losses [Cooper and Knittel, 2006, Altshuler et al., 2009] and traces the asymmetry's effects on investment and risk-taking [Edgerton, 2010, Dreßler and Overesch, 2013, Ljungqvist et al., 2017, Langenmayr and Lester, 2018, Goodman

et al., 2023, Hillmann and Jacob, 2025]. We add causal estimates of what relaxing the asymmetry does, in administrative data covering the private and small firms where losses concentrate, and on margins, survival and employment, that this literature has rarely been able to observe.

The second literature concerns the refunding of losses specifically. Zwick [2021] shows that most eligible American firms leave refunds unclaimed and traces the gap to tax preparation frictions; Dobridge [2021] and Bethmann et al. [2018] study how refund dollars are used, finding effects concentrated among constrained firms; Koch et al. [2023] find that generous loss offsets stabilize equity values in crises; and Langenmayr and Lester [2018] link longer refund windows to greater risk-taking. Relative to this work, we offer a research design that does not compare claimants to non-claimants or rely on cross-country variation, population coverage that includes the private firms most existing samples miss, and outcomes that extend to exit and employment. Our take-up results also carry the American puzzle to a new setting: even in a crisis, with no cap on the refund, most eligible firms left the money unclaimed.

The third literature evaluates support to firms and workers during crises: short-time work schemes that pay firms to keep workers on reduced hours [Giupponi and Landais, 2023, Cahuc et al., 2021], guaranteed lending and wage subsidies during the pandemic [Autor et al., 2022, Granja et al., 2022, Gourinchas et al., 2025], and investment stimulus through the tax code [Zwick and Mahon, 2017]. The instrument we study differs from these in ways that matter for policy design. Its targeting is automatic, conditioning on the firm's own tax history rather than on applications or discretion, which makes it hard to manipulate and directs larger refunds to firms with longer records of paying tax. Much of its outlay is recovered, because refunds are advances against deductions the recipients would otherwise have taken. And it requires no new administrative apparatus, operating entirely through the existing tax return. Our estimates make it possible to set this instrument alongside these alternatives.

Section 2 sets out a framework for valuing refunds to firms and costing them to the government. Sections 3 and 4 describe the Australian setting and the data, and Sections 5 through 8 present the empirical work, and Section 9 concludes.

2 Conceptual Framework

This section sets out the economics of introducing a loss carryback into a corporate tax system that already permits losses to be carried forward. Here we describe a generic asymmetric corporate income tax rather than any particular jurisdiction's, and defer the institutional detail of the Australian system to Section 3. A carryback does not create new loss relief; it changes the timing and the certainty of relief the system already provides. Section 2.1 asks what that change is worth to the firm, and Section 2.2 asks what it costs the government.

2.1 Firm gains from carrybacks

Consider a corporate income tax levied at a constant proportional rate $\tau \in (0, 1)$ on a firm's taxable profit. The tax is asymmetric: positive profits are taxed, but negative profits generate no payment from the government. Taken in isolation, a period with pre-tax profit π generates a liability of $\tau \max\{\pi, 0\}$. Loss relief provisions partially undo this asymmetry by allowing a loss in one period to reduce taxable profit in another. We take as the status quo the regime that prevails in Australia and in many other systems: indefinite carryforward and no carryback. A loss may be applied against taxable income in any later period, without expiry, but never against income already reported.[2]

[2]Some countries limit the number of years a loss may be carried forward. The United States, for example, only allows losses to be carried forward for 20 years. We abstract from this as this limitation is not binding for the

A firm's taxable profit is not its accounting profit: taxable income is computed under statutory rules that depart from financial reporting standards in the treatment of depreciation, capital expenditure, and the timing at which revenue and expenses are recognized. A period in which allowable deductions exceed assessable income generates a tax loss, or in the United States a net operating loss; we use the two terms interchangeably. Tax losses are common, reported by 40 to 50 percent of firms in a given year, and they are only loosely connected to accounting losses or to economic distress: accelerated depreciation, immediate expensing, and deductions carried forward from earlier years can each drive taxable income below zero at a firm whose operations are sound.

Under this regime, a tax loss is an asset with two unattractive properties [Favilukis et al., 2025]. It is illiquid, in that it can be realized only by earning future taxable income, and cannot be sold, transferred outside the firm, or straightforwardly borrowed against. And it is contingent, in that its value is realized if and only if such income materializes; a firm that never returns to profit never realizes it at all. The asset also bears no interest: a dollar of loss shelters exactly one dollar of future income however long it is held. A carryback converts this asset into cash. The firm applies the loss against profits it has already reported and receives a refund of tax it has already remitted.

To make the comparison precise, index periods by $t \in \{-1, 0, 1\}$. The firm earns $\pi_{-1} > 0$ in the first period and remits $\tau\pi_{-1}$. In the second it incurs a loss, written $L_0 \equiv -\pi_0 > 0$. Third-period profit π_1 is stochastic, with a distribution that may depend on the firm's taxable profit history. The reform under study permits the firm, at $t = 0$, to apply its loss against period -1 income and claim a refund of

$$R = \tau \min\{L_0,\, \pi_{-1}\}. \tag{1}$$

The refund is capped twice: by the size of the loss, and by the prior year's taxable income, or equivalently, by the tax the firm has already remitted and that can therefore be returned to it. Valuing the reform requires valuing the alternative the firm gives up. A firm that does not carry its loss back carries it forward and deducts it against future taxable income. Let $\beta \in (0, 1)$ denote the firm's discount factor, and suppose for the moment that the firm is risk-neutral and financially unconstrained. Let p denote the probability that the firm subsequently generates enough taxable income to absorb the loss in full; the parameter compresses whether the firm survives to earn taxable income at all, and whether that income suffices to exhaust the loss.[3] Under carryforward the firm obtains a deduction worth τL_0 with probability p and nothing otherwise, one period hence, so the value of the carryforward is

$$V^{cf} = \beta p\, \tau L_0. \tag{2}$$

Carrying the loss back instead delivers τL_0 immediately and with certainty. We take $L_0 \leq \pi_{-1}$ for exposition, so that the refund is not capped. The value of the carryback relative to the carryforward is therefore

$$V^{cb} = \tau L_0 - \beta p\, \tau L_0 = \tau L_0\,(1 - \beta p) = \underbrace{\tau L_0\,(1-\beta)}_{\text{present-value gain}} + \underbrace{\tau L_0\,\beta\,(1-p)}_{\text{certainty gain}}. \tag{3}$$

Equation (3) separates two distinct reasons why an immediate refund is worth more than a deferred deduction of the same nominal size. The first term is a pure timing gain: even were the

majority of firms. Countries that do not tax retained or reinvested earnings, such as Estonia and Georgia, do not permit carrying forward of losses at all.

[3] The general case replaces βp by the discounted utilization rate $\phi_\beta \equiv \sum_{s \geq 1} \beta^s u_s$, where u_s is the share of the loss deducted s periods after it is incurred, and every expression below goes through unchanged.

deduction certain to be taken, the firm would rather have the money now, and the gain is its own discount applied to the full value of the loss. The second term is a contingency gain: with probability $1 - p$ the deduction is never taken at all, and its value is lost. Under risk neutrality this second term is the expected value of deductions that would have gone unused, rather than a premium for bearing risk.

Expression (3) prices the carryforward at its expected value, and for that reason understates the gain from a carryback. The state of the world in which a carryforward is relieved is the state in which the firm earns taxable income, which is also the state in which a marginal dollar is worth least to it. Relief is procyclical, and a claim that pays only in good states is worth less than a claim paying the same amount on average but paying it evenly. The government's claim on the firm is in this sense an option written on pre-tax income: what a deduction is worth depends on the states in which it is taken and not only on how often [Green and Talmor, 1985, Favilukis et al., 2025].

Let m denote the firm's stochastic discount factor between the loss year and the year in which relief would arrive, normalized so that $\beta = \mathbb{E}[m]$, and let $\mathbf{1}$ be the indicator that the loss is relieved. The carryforward is then worth

$$V^{cf} = \mathbb{E}[m\mathbf{1}]\,\tau L_0 \equiv \beta q\,\tau L_0, \qquad q \equiv \frac{\mathbb{E}[m\mathbf{1}]}{\mathbb{E}[m]}, \tag{4}$$

where q is a risk-adjusted utilization rate: the probability that the loss is relieved, reweighted by the value to the firm of a dollar received in the state in which relief arrives. In the two-state case it is simply the risk-neutral probability that the firm returns to profit. Because relief coincides with profitability, $\mathrm{Cov}(m, \mathbf{1}) \leq 0$ and therefore $q \leq p$, with equality only if the timing of relief is independent of the firm's marginal valuation of funds.

Substituting βq for βp in (3) gives

$$V^{cb} = \underbrace{\tau L_0\,(1-\beta)}_{\text{present-value gain}} + \underbrace{\tau L_0\,\beta\,(1-p)}_{\text{certainty gain}} + \underbrace{\tau L_0\,\beta\,(p-q)}_{\text{risk premium}}. \tag{5}$$

The third term is what the firm would pay to exchange a procyclical claim for cash. It vanishes under risk neutrality, in which case (5) collapses to (3), and it is largest for firms whose return to profit is most exposed to aggregate conditions, and hence largest in precisely the episodes in which carrybacks are enacted. Section 2.2 shows that the government does not value the same forgone deduction at q, and that the difference between the two valuations is a gain from trade rather than a transfer.

The reform enlarges the firm's choice set without withdrawing anything from it, which is why carrybacks are conventionally regarded as a unilateral gain to firms. Because $\beta \leq 1$ and $p \leq 1$, expression (3) is non-negative, and strictly positive unless the firm neither discounts the future nor faces any prospect of failing to use the loss; and because the election is voluntary, a firm claims only when it prefers the refund to the deduction. Three assumptions are imposed in delivering this conclusion. First, the tax rate is constant. If the rate is expected to rise between the loss year and the year in which the deduction would have been taken, a deferred deduction is worth more per dollar of loss, and the carryback dominates only if $\tau_0 \geq \beta p\,\mathbb{E}[\tau_1]$. Second, claiming is costless. In practice, claiming a carryback refund requires an active election and perhaps the professional advice to identify it, and Zwick [2021] documents that small firms leave substantial refunds unclaimed; we return to take-up in Section 5. Third, the refund is unconstrained beyond (1). Real systems impose further caps, and Australia's dividend imputation system limits the refund to the firm's stock of corporate tax not yet passed through to shareholders, as described in Section 3.

Expression (3) also holds p fixed. It treats the firm as continuing to operate whatever happens at $t = 0$, so that the refund alters only the timing and certainty of a payment the firm was going to

receive in any case. But the refund arrives as cash precisely in the period in which the firm is losing money, and a substantial literature identifies illiquidity, rather than fundamental insolvency, as a proximate cause of firm exit [Schmalz et al., 2017, Gourinchas et al., 2025]. Suppose the firm must end period 0 with cash balances at or above a floor $c \leq 0$, reflecting its reserves and its borrowing capacity, and ceases operating otherwise, so that it survives if and only if $\pi_0 + R \geq c$. For a firm comfortably above this threshold the refund does nothing to survival and (3) is the whole of its value. For a firm below it, the refund is the difference between continuing and closing, and its value is not (3) but the refund plus the entire continuation value of the enterprise,

$$V^{cb} = R + \beta v, \qquad v \equiv \mathbb{E}\left[\pi_1 - \tau \max\left\{\pi_1, 0\right\}\right], \tag{6}$$

where v is expected post-tax profit in the following period.

The distinction implies, first, that p is itself an outcome of the policy rather than a primitive, since a firm that fails cannot use a carryforward: the probability that the loss would have been relieved in the counterfactual is endogenous to whether the refund is paid. Second, it implies that the value of a carryback is highly heterogeneous across firms holding identical losses, and that the relevant scaling is the size of the refund relative to the firm's cash position rather than the size of the refund alone. Two firms with the same entitlement can be in entirely different positions if one has ample reserves and the other none. Third, as the next subsection shows, this heterogeneity is also what determines the cost of the reform to the government.

2.2 The fiscal cost of carrybacks

Consider now the government's side of the same carryback refund. Let β_g denote the government's discount factor. In every case the government pays out R at $t = 0$. What differs across firms is what it gets back, and the partition induced by the liquidity constraint in Section 2.1 is exactly the partition that organizes the fiscal accounting. Table 1 collects the three cases.

Table 1: Carrybacks as transfers, loans, and lifelines

Firm type	Survival condition	Firm's gain	Government's net cost	Character
Inframarginal exit	$\pi_0 + R < c$	R	R	Transfer
Inframarginal survivor	$\pi_0 \geq c$	$R(1 - \beta q)$	$R(1 - \beta_g p)$	Loan
Marginal survivor	$\pi_0 < c \leq \pi_0 + R$	$R + \beta v$	$R - \beta_g \tau \mathbb{E}\left[\pi_1^+\right]$	Lifeline

Notes: The table records the gain to a firm receiving a carryback refund R and the corresponding net cost to the government, for each of the three firm types induced by the liquidity constraint. τ is the corporate tax rate, β and β_g the firm's and the government's discount factors, p the probability the loss would otherwise have been absorbed by future taxable income, $q \leq p$ its risk-adjusted counterpart defined in equation 4, c the firm's cash floor, $v \equiv \mathbb{E}\left[\pi_1 - \tau\pi_1^+\right]$ its expected post-tax continuation value, and $\pi_1^+ \equiv \max\{\pi_1, 0\}$. All quantities are measured relative to the status quo of indefinite carryforward with no carryback. Firm gains are derived in Section 2.1 and government costs in Section 2.2.

Consider first a firm that would have survived and returned to profit regardless of the refund, which we term an inframarginal survivor. Such a firm does not acquire a deduction; it exchanges one. The refund it receives today is precisely the relief it would otherwise have obtained later, so the government's future collections rise by the amount of the deduction forgone. Discounting at β_g, the government recovers $\beta_g p R$ in expectation, and its net cost is

$$NC = R(1 - \beta_g p). \tag{7}$$

The outlay is not a grant but an advance against the firm's own tax asset: the government pays now and is repaid through higher collections later. Were the loss certain to be used and the government not to discount, the net cost would be zero. The true fiscal cost is the interest on an interest-free loan, and it is small in proportion to the headline outlay.

This wedge is not the one that appears in the firm's valuation. The firm gives up $\beta q R$ and the government expects to recover $\beta_g p R$ on the same forgone deduction, and the two differ in the probability as well as in the rate. The valuations differ because the government holds a large portfolio of such claims and diversifies across firms and over time the risk that any particular loss goes unused; a single firm cannot, and values the same claim at q rather than at p. Setting $\beta = \beta_g$ to isolate this channel, the ratio of what an inframarginal survivor gains to what the refund costs the government is

$$\frac{R\left(1-\beta q\right)}{R\left(1-\beta_g p\right)} = \frac{1-\beta q}{1-\beta_g p} \geq 1, \tag{8}$$

with strict inequality whenever relief is procyclical. The marginal value of public funds exceeds one for these firms even though no real activity changes and the refund is, mechanically, an interest-free advance against an asset the firm already holds. What the reform provides them is not liquidity but insurance, and the surplus arises because the government is the cheaper bearer of the risk that a loss is never relieved. This is the loss-offset counterpart of the risk-sharing role of the corporate tax documented by Ljungqvist et al. [2017]. It also implies that the conventional case for carrybacks, which rests on the marginal survivors in the third row of Table 1, understates the gain from the reform: the second row, which is much the largest, is not welfare-neutral.

At the opposite extreme is a firm that fails whether or not it receives a refund, an inframarginal exit. Such a firm would never have generated taxable income against which to deduct its loss, so there is nothing for the government to recover. The full R leaves the corporate tax base permanently, as does the loss that purchased it, and the net cost is R. The refund is a pure transfer to the firm's claimants. Whether such a transfer is objectionable depends on who ultimately receives it: creditors, shareholders, or, indirectly, workers. This is a question about incidence rather than about fiscal cost.

Between these extremes lie the firms whose survival the refund changes. Here the government pays R and acquires something it would not otherwise have had: a firm that continues to operate, to remit tax, and to employ workers. Writing $\pi_1^+ \equiv \max\left\{\pi_1, 0\right\}$, its net cost is

$$NC = R - \beta_g \tau \, \mathbb{E}\left[\pi_1^+\right], \tag{9}$$

the outlay less the discounted tax collected from a firm that would otherwise have ceased to exist. This term can be substantial, and in principle can exceed the outlay. It is also the only one of the three cases in which the government's payment purchases a change in real activity rather than a change in the timing of a payment.

The three cases differ in a single parameter. Define the utilization rate ϕ as the share of the loss the firm would eventually have deducted in the absence of the carryback, so that $\phi = p$ in the all-or-nothing case above and $\phi = 0$ for any firm that exits. The net fiscal cost of a dollar of refunds is then $1 - \beta_g \phi$, and ϕ alone locates a given refund on the spectrum between a loan and a grant. Where ϕ is close to one the government is lending against an asset the firm already holds, and the outlay is largely recovered; where ϕ is close to zero it is giving, and the outlay is not recovered at all. Yet ϕ is not a policy parameter. The government chooses who is eligible and for how much, but not the rate at which the losses it refunds would otherwise have been used. That rate is an empirical object, determined jointly by which firms hold losses and by whether those firms return to profit.

This is why the eligibility criterion matters more than its mechanical simplicity suggests. Carryback eligibility conditions on prior tax paid, which is to say on a history of profitability. In so far as past profitability predicts future profitability, the criterion tags firms with high ϕ, and the outlay behaves like a loan; in so far as it does not, because a firm's losses reflect permanent decline rather than a transitory shock, the same criterion delivers refunds to firms that will never pay tax again, and the outlay behaves like a grant. The criterion has the further attraction of being difficult to manipulate, since it depends on tax remitted in a year that has already closed. Whether it tags well is a question about the joint distribution of past and future profitability, and we take this to the data in Section 5.

Introducing carrybacks therefore does more than accelerate a payment the firm was owed: it shifts risk from the firm to the government. The government charges no interest on the advance, so the discount wedge $\beta_g - \beta$ accrues entirely to the firm. And repayment is contingent: the government is repaid only if the firm returns to profit, and receives nothing if it does not. When carrybacks are enacted during a crisis, the risk is shifted in precisely the states of the world in which the firm's own cost of capital is highest and its capacity to bear risk lowest.

3 Institutional Setting

3.1 Corporate taxation in Australia

Australian businesses operate under one of four legal forms, and the distinction between them governs how their income is taxed. A company is taxed as a separate entity: it pays corporate tax on profits as they accrue, and its shareholders are taxed only when profits are distributed.[4] The remaining three forms are pass-throughs, which face no entity-level tax. A partnership divides its income among the partners, a trust distributes it to beneficiaries, and a sole trader is taxed on business income as personal income; in each case the income is assessed in the hands of the owner at the owner's marginal rate.

Companies pay tax at a single proportional rate on taxable profit; the rate a company faces depends on its size. The standard rate is 30 percent. A reduced rate applies to base rate entities, defined as companies whose aggregated turnover falls below a threshold and no more than 80 percent of whose assessable income is passive.[5] The threshold has been AUD $50 million since 2018-19, and the reduced rate has fallen over time: 27.5 percent through 2019-20, 26 percent in 2020-21, and 25 percent from 2021-22 onward. Before 2017-18 the reduced rate was tied to small business entity status and a lower turnover threshold, and before 2015-16 a single 30 percent rate applied to all companies. Table A.5 in the appendix reproduces the full schedule. Pass-through income is instead assessed under the progressive personal schedule. Tax losses may be carried forward indefinitely, subject to continuity of ownership or a similar business test, but Australia has no general provision allowing them to be carried back.

One further feature of the Australian system matters for the policy we study because it caps the carryback refund: the integration of corporate and personal taxation through dividend imputation, known as franking. When a company pays corporate tax, the payment is credited to a franking account: a running ledger of tax paid at the entity level that has not yet been passed through to shareholders. When the company pays a franked dividend, it attaches franking credits and debits the account by the same amount; the shareholder applies the credits against personal tax, with any excess refunded. Distributed profits are therefore taxed once, at the shareholder's marginal rate,

[4] This is comparable to a C Corp in the United States.

[5] Passive income includes dividends and attached franking credits, royalties, rent, interest, and net capital gains.

and corporate tax operates as a withholding tax on distributions.

3.2 The introduction of carrybacks in Australia

Australia has never maintained a permanent carryback provision. Loss-making companies may carry losses forward indefinitely, but ordinarily cannot apply them against tax already paid. This contrasts with the United States, which has had a carryback of some duration for most of the past half-century and has extended it in every recession since 2001. Australia has instead introduced carrybacks twice, in each case temporarily and in response to a crisis. The first, following the global financial crisis, capped the eligible loss at AUD $1 million, operated for the single income year 2012-13, and was repealed in 2014.[6]

The second episode is the subject of this paper. The temporary loss carry-back was announced on 6 October 2020 as part of the JobMaker Plan in the 2020-21 Budget and legislated within days by the Treasury Laws Amendment (A Tax Plan for the COVID-19 Economic Recovery) Act 2020. Companies could carry losses incurred in 2019-20, 2020-21 and 2021-22 back against tax paid in 2018-19 or any later year; the 2021-22 Budget extended the measure by one year to cover 2022-23 losses. Unlike the 2012-13 provision there was no monetary cap on the loss or the refund. Treasury costed the measure at approximately AUD $4.9 billion over the forward estimates, and designed it to work alongside temporary full expensing, which by accelerating depreciation deductions generates precisely the losses a carryback converts into cash.

Eligibility was defined by legal form, size, and tax history. The claimant had to be a corporate tax entity throughout the loss year, the year the loss was carried back to, and every year in between; partnerships, trusts and sole traders were statutorily excluded regardless of the losses they sustained. Aggregated turnover in the loss year had to be below AUD $5 billion, a threshold that excluded only a handful of the largest firms in the country. The firm had to have been liable for income tax in the year it carried the loss back to, and to have lodged returns for the claim year and the five preceding years, or to have been assessed or not required to lodge in those years. Only revenue losses qualified: capital losses, losses arising from the conversion of excess franking offsets, and losses transferred in by head companies of consolidated groups or by foreign banking groups were all ineligible.

Two design features shape the empirical work that follows. First, the benefit took the form of a refundable tax offset claimed in the return for 2020-21, 2021-22 or 2022-23. A firm with a 2019-20 loss therefore could not monetize it until it lodged its 2020-21 return, which delayed relief for firms hit earliest in the pandemic. The offset equals the loss carried back, reduced by the net exempt income of the year it is carried back to, valued at the corporate tax rate of the loss year; it is capped at that earlier year's income tax liability and, in total, at the firm's franking account surplus at the end of the claim year. Second, the claim was elective. There was no obligation to claim, and no ordering rule requiring losses to be applied to the earliest available year, so an eligible firm could decline the refund and preserve the loss as a carryforward. Take-up is therefore an outcome to be explained rather than a mechanical consequence of eligibility. Integrity provisions denied the offset where control of the company changed through share disposals entered into for the purpose of obtaining it.

As described in Section 3.1, a company's franking account is credited when the company remits

[6] Announced in the 2012-13 Budget and legislated as Division 160 of the Income Tax Assessment Act 1997, the provision was intended to become a permanent two-year carryback, but a one-year transitional period applied in its first year, so that losses incurred in 2012-13 could be carried back only to 2011-12. The cap implied a maximum refund of AUD $300,000 at the 30 percent rate then in force, and the refund was further capped at the prior year's income tax liability and at the firm's franking account balance. The provision was repealed alongside the minerals resource rent tax, with effect from 2013-14.

corporate tax and debited when it attaches franking credits to a distribution, by $\tau/(1-\tau)$ per dollar of franked dividend paid. The closing balance is therefore the stock of corporate tax the company has paid and not yet passed through to its shareholders. The legislation requires that the loss carryback offset not exceed this balance at the end of the year in which the claim is made. Writing F for the closing franking account surplus and $\tilde{R}$ for the offset the firm would otherwise be entitled to, after applying the loss-year rate and the gain-year liability limits, the refund the firm can actually obtain is $\min\left\{\tilde{R}, F\right\}$. The rationale is that a refund of corporate tax must be matched by a reversal of the credits that tax generated: absent the cap, a company could distribute franking credits to its shareholders and then reclaim the underlying tax from the government, so that the same dollar of corporate tax was relieved twice.

3.3 Other support for firms during the pandemic

The carryback was one component of a much larger package of business support. The largest measure was JobKeeper, a wage subsidy paid fortnightly to employers whose turnover had fallen by a prescribed threshold, which at its peak supported roughly 3.5 million workers at a cost near AUD \$89 billion. Other measures included the cash-flow boost, which refunded amounts withheld from employee wages to smaller employers; temporary full expensing, which allowed immediate deduction of eligible capital expenditure; the SME loan guarantee scheme; the JobMaker hiring credit; and various state-level payroll tax deferrals and relief packages.

4 Data

4.1 The Business Longitudinal Analysis Data Environment

Source and coverage Our primary data source is the Business Longitudinal Analysis Data Environment, an integrated administrative database maintained by the Australian Bureau of Statistics (ABS). It is assembled by linking the national business register, the administrative list of every entity holding an Australian business number, to the records the Australian Taxation Office collects in the ordinary course of administering the tax system, and to the Bureau's own business surveys. Coverage is universal: every business that lodges a tax return, remits an activity statement, or reports payroll appears in the data. The panel begins in the 2001-02 financial year and, in its most recent release, extends through 2024-25, although the income tax records that are central to our analysis are currently available through 2022-23.

Unit of observation The database records information at two levels of aggregation. The first is the tax-filing entity, which is the unit that lodges a return and remits tax. The second is the economic unit, which groups together the filing entities operated by a single enterprise. The distinction matters because larger businesses commonly register several business numbers, one for each operating division, property, or franchised outlet, and lodge a separate return for each. Since the hiring, investment and exit decisions we study are made by the enterprise rather than by the filing entity, we build the panel from filing entities upward and then aggregate to the economic unit, using a concordance ABS supplies for each year. We refer to the resulting aggregated unit as the firm.

Modules The data are organized into modules, of which we use six. The business register module supplies each entity's industry, legal form, institutional sector and state of operation. The activity statement module contains the quarterly returns lodged by every business registered for the goods and services tax, including total sales and wages paid. The income tax module contains the annual

returns themselves, with separate forms for companies, individuals, partnerships and trusts; this is our principal source. The withholding module gives an annual headcount and full-time-equivalent employment count for every employer, back to 2001-02. From 2019-20 onward, the payroll reporting module reports employment, gross pay and tax withheld monthly and in near real time, since employers now report each pay run to the tax authority as it occurs. Finally, a business birth date module records the year in which each entity commenced operating, from which we construct firm age.

The company tax return The company tax return data reports a complete income statement, with revenue and costs itemized and accounting profit reported directly; an abbreviated balance sheet, including current and total assets and liabilities, total debt, and shareholders' funds; and the full reconciliation from accounting profit to taxable income, including exempt income, losses deducted in the current year, and losses carried forward to later years. The last of these is a stock rather than a flow: it is the balance of unused losses the firm carries into the following year, so its lagged value records how much unused relief a firm had already accumulated at the moment it incurred a new loss. We use that balance as a state variable throughout, because the framework of Section 2 implies that the value of a refund to the firm and the rate at which the government recovers it both depend on it, and neither is determined by the size of the current loss alone. And it reports the entire tax calculation, so that we observe not merely taxable income but the tax charged on it, the offsets applied against that charge, and the net amount due or refundable at lodgement. The returns lodged by pass-through entities are considerably sparser, containing income, expenses and net income on a comparable basis.

Two further groups of items on the company return are central to this paper. The first concerns dividend imputation. Because Australia integrates corporate and personal taxation through the franking system described in Section 3, the return reports franked and unfranked dividends paid, franking credits received on dividends from other companies, and the opening and closing balance of the firm's franking account -- the running record of corporate tax paid that has not yet been passed through to shareholders. As we explain there, this balance caps the refund a firm can obtain from a carryback, so observing it directly is essential rather than incidental. The second group was created specifically for the carryback policy and appears only on returns lodged for the 2020-21 through 2022-23 income years. For each admissible pairing of a loss year with an earlier profitable year, the return records the amount of loss the firm elected to carry back. It separately records the corporate tax rate applying in each loss year and the net exempt income and income tax liability of each earlier year, which together are the inputs to the statutory refund formula, and it reports the resulting refundable offset within the tax calculation. We therefore observe claims exactly, and at the level of the individual loss-year-to-profit-year pairing, for every company that made one.

Constructed variables Two quantities that the analysis requires are not reported and must be constructed. The first is the statutory tax rate each firm faced in each year. Australia's corporate rate has varied over our sample both across time and across firms, since companies below a turnover threshold pay a reduced rate. We recover the rate a firm actually faced as the ratio of the tax charged on its taxable income to that taxable income. Where this ratio cannot be computed, or falls outside the range of rates in force that year, as it does for firms without positive taxable income and for the small number of entities taxed under special schedules, we impute it from the published rate schedule using the firm's reported small-business and reduced-rate status together with its aggregated turnover. The second constructed quantity is the refund a firm could have obtained from a carryback had it chosen to claim one, which we discuss in the next subsection. All dollar

amounts throughout the paper are deflated to 2023 Australian dollars. Appendix A documents every step of the construction and reports the source variable behind each measure used in the paper.

4.2 Analysis sample

We begin from the universe of economic units and the tax-filing entities that compose them. We drop filing entities that the Bureau's concordance assigns to more than one economic unit in a given year, since their reported tax and payroll aggregates cannot be attributed cleanly to a single enterprise. We drop non-profit institutions and entities in the government sector. The result is a panel indexed by filing entity and financial year, into which we merge the income tax returns, the activity statements, the annual and monthly payroll records, and the business birth dates. We then aggregate to the firm. Flows and stocks are summed across a firm's constituent filing entities. Categorical attributes such as industry, legal form and state are taken from the largest constituent by income, and ratios are recomputed from the aggregates rather than averaged across constituents.

Two restrictions define the estimation sample. We retain firms that employ at least five people in two consecutive years and that report more than AUD \$150,000 of revenue, in 2023 dollars, in two consecutive years. The business register contains an enormous number of entities that are registered but essentially dormant: shell companies, holding vehicles, single-property trusts, and one-person service entities whose payroll and tax records are mechanical artifacts rather than measurements of real activity. Because our outcomes are employment and firm survival, retaining these entities would add noise without adding information, and would make the corporate and pass-through populations difficult to compare on any margin we care about. We impose the conditions over two consecutive years rather than a single year so as not to select on the transitory shocks that are themselves the object of study.

The central variable in the analysis is the refund a firm was entitled to claim, whether or not it in fact claimed one. This is not observed. The carryback items on the tax return are populated only for firms that lodged a claim; for a firm that was entitled to a refund and left it unclaimed the data records a missing value for the carryback variables. We therefore construct each firm's potential carryback by applying the statutory formula to its own reported history. For a firm with a loss in 2020-21, the loss is applied first against taxable income in 2019-20 and then, if any remains, against 2018-19, with the amount available at each step reduced by that year's exempt income. The loss so applied is valued at the corporate tax rate of the loss year, and the resulting offset is capped at the income tax liability of the year against which it is applied and, in aggregate, at the balance of the firm's franking account at the end of the claim year. Appendix A.7 sets out the calculation for each claim year together with the variables used at each step.

Because this measure is constructed rather than observed, its quality governs everything that follows. Among firms that did claim a carryback, regressing our constructed potential carryback on the refund actually claimed yields an R^2 of 0.977, with a slope of 0.988 and a standard error of 0.164. Two things follow. The first is behavioral: conditional on claiming at all, firms take essentially the entire refund available to them, so that the amount claimed and the amount available coincide. The second is that our reconstruction of the statutory entitlement reproduces the entitlement itself almost exactly. That gives us confidence in the measure for the far larger population of eligible firms that never claimed, and for which no direct benchmark exists.

The analysis sample retains firms with strictly positive potential carryback in 2020-21. This leaves 56,330 firms. Section 5 reports summary statistics for this sample alongside the broader populations from which it is drawn.

5 Facts about carrybacks in Australia

This section documents who the carryback reached before the remainder of the paper asks what it caused. It establishes three sets of facts: which firms the eligibility criterion selects and how those firms subsequently fare; how carryback capacity is distributed across firms; and which eligible firms claimed. The targeting facts measure the tagging content of the eligibility criterion, the distribution facts locate the firms for which a refund is large relative to the scale of their operations, and the take-up facts reveal which firms found the refund worth claiming. Throughout, the unit of observation is the firm in the estimation sample of Section 4, and the loss year is 2020-21 unless stated otherwise.

5.1 What the eligibility criterion selects

Eligibility combined a state and a history. A firm had to sustain a tax loss in 2020-21; it had to have paid income tax in a year the loss could be carried back to; it had to hold a positive franking balance at the end of the claim year; and its aggregated turnover had to fall below AUD $5 billion. The state was widespread: 31 percent of firms in the estimation sample report negative taxable income in 2020-21. The history was the binding condition: of these loss-makers, 42 percent meet the remaining requirements and could have claimed a refund. Which firms clear that condition is the first fact of interest, because Section 2 shows that the fiscal character of the policy turns on whom the criterion selects.

Table 2 compares four nested populations in the 2020-21 year: all firms in the estimation sample, firms with negative taxable income, firms eligible for a carryback, and firms that claimed one. Three comparisons summarize the targeting. First, eligible firms are larger than the average loss-maker on every margin: mean revenue of AUD $1.5 million against AUD $1.1 million, 8.3 employees against 7.7, and a wage bill roughly a quarter larger. They also pay more per worker and hold more assets. Second, they enter the loss year from far stronger histories, a fact Figure 1 documents directly: a refund requires prior tax paid, and prior tax paid requires prior profits, so this selection is built into the criterion mechanically. Third, claimants resemble the eligible population from which they are drawn far more than either resembles the average loss-maker, and are if anything somewhat larger. The average refund among claimants is AUD $46,000; averaged over all eligible firms, with non-claimants at zero, refunds average AUD $13,200.

Table 2: Firm characteristics by population, 2020–21

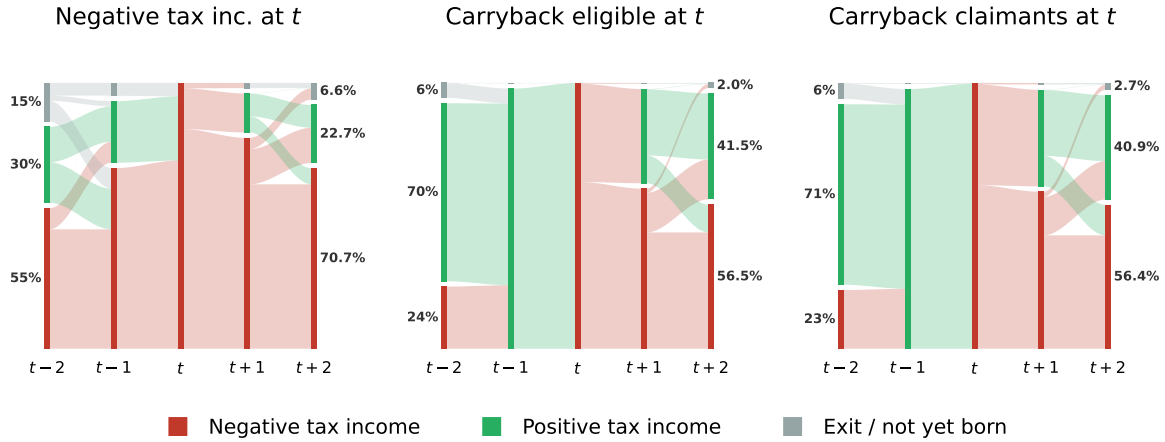
	All firms	Neg. taxable inc.	CB eligible	CB claimants
Firms	438,240	135,080	56,330	16,150
Revenue	1,682,422	1,086,623	1,512,176	1,815,334
Employment	9.4	7.7	8.3	9.6
Wage bill	744,125	424,475	534,348	725,127
Total assets	832,684	648,328	8,362,270	1,058,181
Carryback amount	1,695	5,144	13,184	45,988
Share claiming	0.04	0.11	0.29	1.00

Notes: Columns report means over four nested populations in the 2020–21 income year: all firms in the estimation sample, firms with negative taxable income, firms eligible for a carryback, and firms that claimed one. Rows report firm counts, mean revenue, employment, the wage bill, total assets, the refund claimed, averaged over all firms in the group with non-claimants at zero, and the share of firms claiming. Dollar amounts are deflated to 2023 Australian dollars. Counts and statistics are rounded in accordance with ABS disclosure rules.

The consequences of that selection for firm trajectories are visible in Figure 1, which follows three populations defined in the loss year, all loss-makers, the eligible, and claimants, over the two

years on either side. The panels differ starkly in where firms come from. Two years before the loss, 70 percent of eventually-eligible firms had positive taxable income and 24 percent were making losses; among all loss-makers the corresponding shares are 30 and 55 percent, with the remainder not yet operating. They differ equally in where firms go. Two years after the loss, 41.5 percent of eligible firms have returned to positive taxable income and only 2.0 percent have ceased operating, against 22.7 and 6.6 percent among all loss-makers: eligible firms recover at roughly twice the rate of the average loss-maker and exit at a third of it. Claimants are nearly indistinguishable from the eligible population on both sides of the loss year. The criterion therefore tags exactly the firms the fiscal logic of Section 2 favors, those whose losses are transitory interruptions of otherwise profitable histories. But the same figure is a warning against causal readings of any of these comparisons: firms selected on profitable histories would have recovered faster than the average loss-maker under any policy, and separating the refund's contribution from that selection is the task of Section 6.

Figure 1: Taxable income trajectories around the 2020–21 loss year



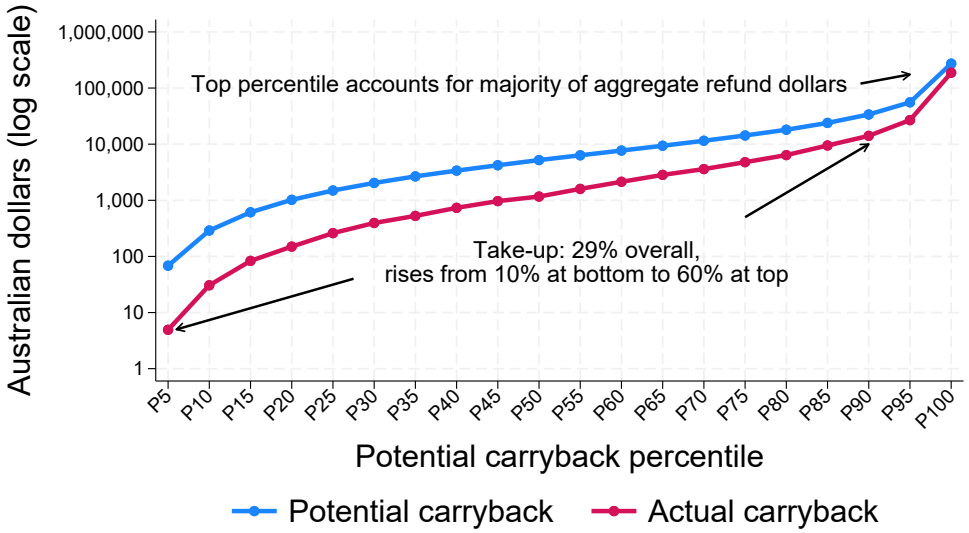
Notes: Each panel follows a fixed population of firms defined in the 2020–21 loss year, denoted t: firms with negative taxable income (left), the subset eligible for a carryback (center), and firms that claimed one (right). Columns give the share of each population with negative taxable income, positive taxable income, or not operating (exited or not yet born) in each year from $t-2$ to $t+2$, with flows between states shaded. Shares are rounded in accordance with ABS disclosure rules.

5.2 The distribution of carryback capacity

Section 4 constructed, for every firm, the refund it was entitled to claim. Figure 2 ranks eligible firms into ventiles of this potential carryback and plots the average potential and actual refund within each, on a log scale. The distribution is extremely right-skewed. The average entitlement is roughly AUD \$70 in the bottom ventile, \$5,000 in the middle of the distribution, and more than \$250,000 in the top ventile, a span of more than three orders of magnitude. Aggregate dollars are correspondingly concentrated: the top percentile of entitlements alone accounts for the majority of aggregate refund dollars. The vertical gap between the two series is take-up, to which we return below.

Figure 3 disaggregates by firm size, measured by the firm's percentile of average revenue over the preceding decade. In levels, the mean potential carryback rises steeply with size, from roughly AUD \$2,500 among the smallest employers to more than \$120,000 in the top 0.1 percent, and the top 1 percent of firms by revenue hold roughly 20 percent of aggregate potential refund dollars. Scaled by the firm's own payroll, the gradient largely disappears: the mean entitlement sits between 4 and 9

Figure 2: The distribution of potential and actual carrybacks



Notes: The sample is firms eligible for a carryback in 2020–21, ranked into ventiles of the potential carryback constructed in Section 4.2. The figure plots the mean potential carryback and the mean refund actually claimed within each ventile, in 2023 Australian dollars on a log scale. Means are taken over all eligible firms in the ventile, so the actual series includes non-claimants at zero. Statistics are rounded in accordance with ABS disclosure rules.

percent of the annual wage bill across the entire size distribution. Refund dollars are concentrated among large firms because losses and prior tax payments are, not because the policy is generous to large firms in proportion to their scale. For a firm of any size, an entitlement of several percent of annual payroll, arriving as cash in the middle of a loss year, is first-order relative to the flows that Section 2 identifies as determining survival.

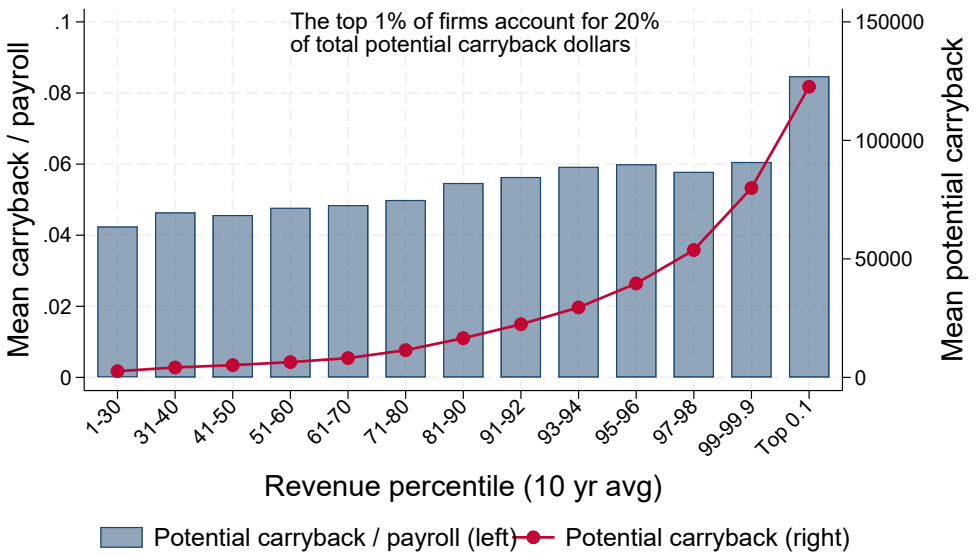
5.3 Take-up

Claiming required an election on the tax return, and most eligible firms did not make one: only 29 percent of eligible firms claimed a carryback in 2020-21. Low take-up of carrybacks is not unique to Australia. Zwick [2021] documents that most eligible United States firms failed to claim refunds they were entitled to, and attributes much of the shortfall to fixed costs of tax sophistication. The Australian pattern is consistent with fixed costs operating here as well. The gap between the potential and actual series in Figure 2 measures this margin: take-up rises from roughly 10 percent in the bottom ventile of the potential carryback distribution to roughly 60 percent at the top. Because entitlements are so skewed, dollar-weighted take-up is considerably higher than the firm-weighted rate: the average refund among claimants is AUD $46,000.

The margin of adjustment is whether to claim, not how much. Among claimants the amount claimed is essentially the whole entitlement, as documented in Section 4, and Appendix Figure B.1 shows the implication: the share of potential refund dollars actually used within each ventile tracks the share of firms claiming almost exactly. The same appendix figure repeats the exercise for Australia's brief 2012-13 carryback, described in Section 3, and finds the same profile, which suggests the pattern is a feature of carryback claiming rather than of the pandemic.

Which firms claim, holding the stakes fixed? Figure 4 plots take-up against the firm's baseline financial position, measured as the ratio of earnings to total liabilities, after residualizing on the size

Figure 3: Carryback capacity by firm size



Notes: Firms are grouped by percentile of average revenue over the ten years preceding the loss year. The line plots the mean potential carryback within each group, in 2023 Australian dollars on the right axis. The bars plot the mean ratio of the potential carryback to the firm's annual payroll, on the left axis. Statistics are rounded in accordance with ABS disclosure rules.

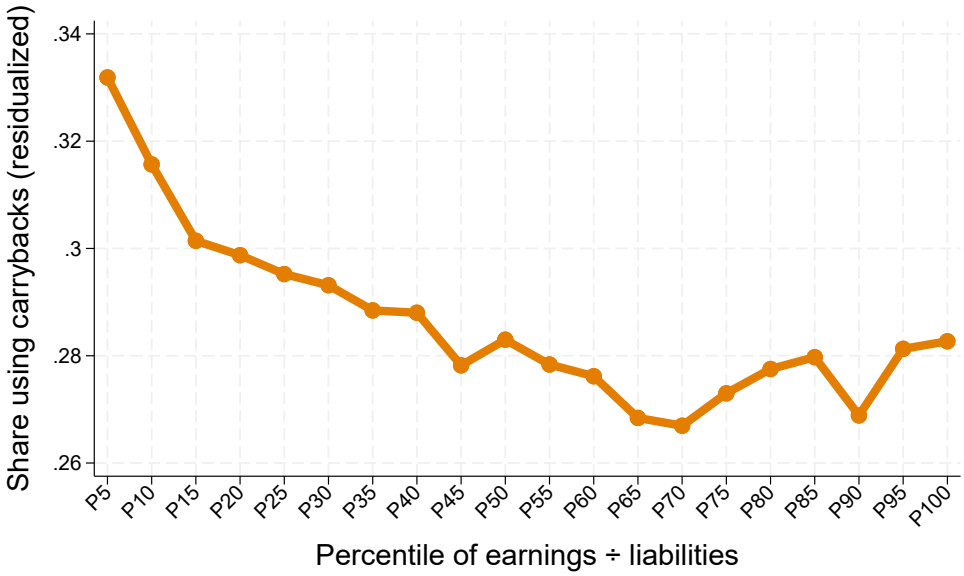
of the potential carryback so that the comparison is between firms with the same entitlement. Take-up declines in financial strength: firms in the bottom ventiles of earnings relative to liabilities claim at rates roughly five percentage points above those at the top. The firms that claim are therefore those for which Section 2 predicts the refund is most valuable: firms whose current earnings least cover their obligations, for which the immediate and certain character of the refund matters most, and for which the refund is most plausibly pivotal to continuing operation. This selection into claiming is informative about the policy's incidence, but it also means that comparisons of claimants with non-claimants confound the refund with the distress that prompted the claim, a point Section 6 takes up directly.

6 Comparing corporations to pass-throughs

The patterns documented in the previous section are informative about who the carryback reached, but not yet about what it did. Eligible firms entered the pandemic larger, older and more consistently profitable than the average loss-maker, and they returned to profitability at roughly twice the rate. Both facts are exactly what the eligibility criterion should produce whether or not the refund itself had any effect. Eligibility conditions on prior tax paid, so the eligible population is selected on the very histories that predict recovery: a firm that remitted tax in 2018-19 differs from one that did not, in its products, its management and its balance sheet, and it would have weathered the pandemic differently even had no refund existed. A comparison of eligible firms with ineligible ones therefore confounds the causal effect of the refund with the selection built into the criterion that assigns it.

Take-up compounds the problem rather than solving it. The claim was elective, and the previous section showed that the decision to claim responded to the firm's circumstances: conditional on the

Figure 4: Take-up and financial fragility



Notes: The sample is firms eligible for a carryback in 2020–21, ranked into ventiles of the ratio of earnings to total liabilities at baseline. The figure plots the share of firms claiming within each ventile, residualized on the size of the potential carryback with the sample mean added back, so that the gradient is net of the relationship between take-up and the entitlement in Figure 2. Statistics are rounded in accordance with ABS disclosure rules.

size of the entitlement, more financially constrained firms claimed more often. If claimants are the more distressed among the eligible, a comparison of claimants with non-claimants understates the benefit of the refund; if fixed costs of claiming deterred the smallest firms, it overstates it. Neither comparison recovers the object of interest, which is the difference between a firm's outcomes with the refund and the outcomes the same firm would have experienced carrying its loss forward. Recovering it requires variation in access to the carryback that is unrelated to a firm's own history and choices. The statutory exclusion of pass-throughs provides that variation, and this section develops the comparison it supports.

6.1 Pass-throughs as a comparison group

As Section 3 described, the carryback was confined to corporate tax entities: partnerships, trusts and sole traders were excluded by statute, whatever losses they sustained and whatever tax their owners had paid. The exclusion follows from the mechanics of the instrument, since a carryback refunds entity-level tax and a pass-through remits none. Its consequence is that the reform changed the tax treatment of losses for one of two legal populations operating side by side in every industry, and left the other unchanged. Pass-through losses continued to be relieved by the routes that had always applied to them: a partnership's tax loss is allocated to its partners in the year it is incurred and deducted against their other income, subject to the non-commercial loss rules where the partner is an individual; a trust's loss cannot be distributed and is carried forward within the trust; a sole trader's loss is deducted against other personal income or carried forward under the same non-commercial loss rules. Each of these provisions predates the pandemic and none was altered during it. The pass-through regime is not identical to the corporate one, and we do not require it to be: a partnership or sole trader loss can offset the owners' other income at once, a partial symmetry that corporations lack. What matters for the design is not that the two regimes coincide but that only

one of them changed. A pass-through in 2020-21 faced precisely the loss rules it had always faced, which is what a comparison group must do.

Nor could a firm opt into treatment. Eligibility required the claimant to have been a corporate tax entity throughout the loss year, the year the loss was carried back to, and every year in between, so a pass-through that incorporated after the announcement gained nothing. Treatment status was fixed by a choice of legal form made years earlier, typically at establishment and for reasons of liability, succession and owner-level taxation. That choice is meanwhile not pinned down by industry or scale. Australia's two largest supermarket chains are companies while their principal competitor operates as a limited partnership, and among the country's major law firms some are incorporated and listed while others remain partnerships.[7] Corporations and pass-throughs of similar size sell in the same markets, hire in the same labor markets, and were exposed to the same public health restrictions. And because both lodge returns with the same tax authority and report payroll through the same monthly system, their outcomes are measured identically in our data.

Comparisons across legal form have a history in public finance. Yagan [2015] uses S corporations, which the 2003 United States dividend tax cut left untouched, as the comparison population for the C corporations the cut affected, documenting that the two populations track each other closely outside the reform; Kennedy et al. [2026] use the same comparison to estimate the effects of the 2017 United States corporate rate cut on firms and their workers. The Australian setting has the same structure with the roles reversed: here the tax change applies to the corporate form, and the pass-through forms supply the counterfactual.

The design is a difference-in-differences in exposure intensity. Section 4 constructed, for every corporation, the refund the statute entitled it to, whether or not it was claimed. The next subsection extends the construction to pass-throughs, computing for each one the refund it could have claimed had it been a corporation with the same history. Write D_i for this entitlement, actual or counterfactual, and Corp_i for an indicator that firm i is a corporation. Among corporations, outcomes may covary with D_i for two reasons: because the refund does something, and because the entitlement is a summary of the profit history that generated it. Among pass-throughs only the second channel operates: their entitlement summarizes the same kind of history, but no refund attaches to it. The differential gradient, the coefficient on the interaction $\text{Corp}_i \times D_i$, is our estimate of the causal effect of the carryback in the estimating equations of Sections 7 and 8.

The identifying assumption is that, conditional on the matching cells and controls described below, the outcomes corporations would have experienced without the reform bear the same relationship to the entitlement as the outcomes of pass-throughs. Writing $Y_i(0)$ for the outcome firm i would have realized in a world without the carryback, measured in levels or, in the event studies, as a change from baseline, we require

$$\mathbb{E}\left[Y_i(0) \mid D_i = d, \text{Corp}_i = 1, X_i\right] - \mathbb{E}\left[Y_i(0) \mid D_i = d, \text{Corp}_i = 0, X_i\right] = g(X_i) \qquad \text{for all } d. \quad (10)$$

The gap between legal forms in counterfactual outcomes may be of any size, but must not vary with the entitlement. Corporations may differ from pass-throughs in the level of every outcome we study; those differences are absorbed by the main effect of Corp_i. Firms with large entitlements may differ from firms with small ones, in their histories above all; those differences are absorbed by the main effect of D_i, which the pass-throughs identify. What the assumption rules out is a factor with three properties at once: it differs across legal forms, it covaries with the size of the entitlement,

[7] Woolworths and Coles, the two largest supermarket chains, are listed companies, while ALDI Australia operates as a limited partnership. Among law firms, Slater & Gordon is incorporated and listed on the Australian Securities Exchange, while Herbert Smith Freehills and King & Wood Mallesons are partnerships.

and it affects outcomes during the pandemic period. A factor with only two of the three does no harm.

The substantive threat is easiest to state through an example. Suppose corporations with long profit histories enjoy better access to credit than pass-throughs with the same histories, perhaps because incorporation itself eases lending against the business. High-entitlement corporations would then have outperformed high-entitlement pass-throughs in the downturn even without the reform, and the interaction would attribute the difference to the carryback. Three features of the setting and the design discipline concerns of this kind. First, no other pandemic-era measure conditioned on legal form: JobKeeper, the cash flow boost and temporary full expensing were each available to companies, partnerships, trusts and sole traders alike (Section 3.3), so the carryback is the only component of the support package that loads on the interaction of legal form with prior profitability. Second, the matched sample constructed below aligns the two populations on the observables most likely to covary with both the entitlement and outcomes: industry, size, wages, age, growth and profitability. The assumption is then asked to hold conditionally and on common support, rather than across two raw populations of very different composition. Third, the assumption has testable implications. Corporations and pass-throughs with the same entitlement should track each other in the months before refunds arrive, and an entitlement constructed by the same formula in a year with no carryback should generate no differential at all; the event studies and placebo exercises below implement these tests.

6.2 Counterfactual carrybacks for pass-throughs

The design requires, for every pass-through, the refund it would have been entitled to as a corporation. We construct it by passing the firm's own reported history through the statutory formula of Appendix A.7, treating its business income as the taxable income of a hypothetical company. The partnership, trust and individual returns each report net business income or loss on a basis comparable to the taxable income of the company return; Table A.4 in the appendix maps the corresponding items. A pass-through whose net business income in 2020-21 is negative is assigned a loss of that amount. Its hypothetical tax payments are computed by valuing net business income in 2018-19 and 2019-20 at the corporate rate the firm would have faced in those years, determined from the published schedule and the firm's aggregated turnover in the same way as the imputed rate for corporations in Appendix A.6. The loss is then carried back exactly as for a corporation: applied first against 2019-20 income and then against 2018-19, capped at each step by the hypothetical tax of the gain year, and valued at the loss-year rate.

One element of the corporate calculation has no pass-through analogue. The franking cap limits a corporation's refund to the stock of corporate tax it has paid but not yet passed through to shareholders, and a pass-through has no franking account, because it remits no entity-level tax to record in one. We therefore construct the pass-through entitlement without the cap, which is equivalent to endowing the hypothetical corporation with full retention of its post-tax earnings. Because the cap binds only for corporations that distributed heavily between the gain years and the claim, the omission bears on the comparability of the two constructions rather than on their logic, and we address it symmetrically: our estimates are essentially unchanged when the franking cap is removed from the corporate entitlement as well, so that the treatment variable is built from an identical formula on both sides.

The remaining conditions are applied symmetrically. Pass-throughs with aggregated turnover above AUD $5 billion in the loss year are assigned an entitlement of zero, as are those without the lodgement history the statute demands of corporate claimants. The entitlement is deflated and scaled by baseline payroll in the same way on both sides, so that D_i carries the same interpretation

for a corporation and for a pass-through: a refund per dollar of wage bill. And like the corporate potential carryback, the counterfactual entitlement is a function of reported history alone. It embeds no decision taken during the pandemic, and in particular no claiming behavior, which pass-throughs were never in a position to exhibit.

6.3 A matched sample of corporations and pass-throughs

The corporate and pass-through populations differ in composition. Pass-throughs dominate among the smallest employers, sole traders and single-outlet trusts above all, while the largest firms are overwhelmingly incorporated. Estimating the differential gradient on the raw populations would therefore rest in part on comparisons between regions of the covariate space where one form or the other is thin. We instead conduct the analysis on a matched sample built to balance the two populations at baseline. Matching is on characteristics measured in the 2019-20 tax year, which closed in June 2020, several months before the carryback was announced, so that baseline characteristics cannot be affected by the policy.

We first estimate a logit for the probability that a firm is a corporation given baseline covariates: the logarithms of revenue, the wage bill and employment; the profit margin; firm age; revenue growth; and taxable income, measured for pass-throughs by net business income. Each corporation is then paired with the pass-through nearest to it in the fitted probability $\hat{p}(X_i)$ within an exact-match cell $c(i)$ defined by two-digit ANZSIC industry, decile of average wage and decile of revenue:

$$m(i) = \underset{j:\,\mathrm{Corp}_j=0,\; c(j)=c(i)}{\arg\min} \left|\hat{p}(X_i) - \hat{p}(X_j)\right|. \tag{11}$$

The exact cells guarantee that every comparison is within industry and size class; the fitted probability orders candidates within a cell by overall similarity. Corporations whose cell contains no pass-through are dropped. The matched sample contains 5,650 corporations and an equal number of pass-throughs.

Matching plays a supporting rather than a leading role in the design. Identification rests on the differential gradient in the entitlement, not on the pairing, and the matched sample serves to enforce common support and to give the conditioning in equation (10) concrete content. Consistent with this, the estimates are similar when the pairing is replaced by entropy reweighting of the pass-through population toward the corporate covariate means; Section 8 reports these checks.

6.4 Comparing the matched populations

Table 3 compares the matched pass-throughs with the corporations they are paired with in the 2019-20 baseline year, reporting group means and tests of their equality. Two features are worth drawing out. First, the two groups are closely aligned on the dimensions the matching targets: revenue, the wage bill, employment, average pay and firm age differ by small and statistically insignificant amounts. Second, the differences that remain are those inherent to legal form itself: total assets differ across the two groups, reflecting the sparser balance sheet reported on the pass-through forms, and we flag this where it bears on the interpretation of magnitudes.

7 The effects of carrybacks on firm survival

The first outcome we examine is whether the firm continues to operate. Section 2 identified survival as the margin on which both the value and the cost of a carryback pivot: for a firm above its cash floor the refund changes the timing of a payment it was owed in any case, while for a firm below it

Table 3: Baseline comparison of matched corporations and pass-throughs

	Corporations	Pass-throughs	p-value of difference = 0
Firms	5,650	5,650	
Revenue	4,782,287	4,782,930	0.997
Employment	17.9	16.7	0.254
Wage bill	1,253,115	1,160,023	0.225
Average wage	67,098	66,346	0.319
Total assets	3,441,912	4,665,644	<0.001
Firm age	15.1	14.7	0.091

Notes: Entries are group means over the matched sample in the 2019-20 baseline year; p-values are from tests of equal means across the two groups. The potential carryback of a pass-through is the counterfactual entitlement constructed in Section 6.2. Dollar amounts are deflated to 2023 Australian dollars. Counts are rounded in accordance with ABS disclosure rules.

the refund is the difference between continuing and closing, and the government's outlay purchases a change in real activity rather than in timing. Which of these cases describes the refunds actually paid is an empirical question, and this section takes it to the design of Section 6.

7.1 Specification

The sample is the matched panel of corporations and pass-throughs constructed in Section 6. The treatment is the entitlement D_i, the potential carryback of the 2020-21 loss year scaled by the firm's baseline payroll, measured for corporations from their own returns and for pass-throughs as the counterfactual entitlement of Section 6.2. A firm is alive at a given horizon if it reports positive payroll in that year, the measure introduced in Section 4. For each horizon $h \in \{1, 2, 3\}$ years after the expected refund we estimate

$$\mathbf{1}\{\text{alive}\}_{i,t+h} = \alpha_h + \gamma_h D_i + \beta_h \,\text{Corp}_i \times D_i + \kappa_{c(i),h} + X_i'\delta_h + \varepsilon_{i,h}, \tag{12}$$

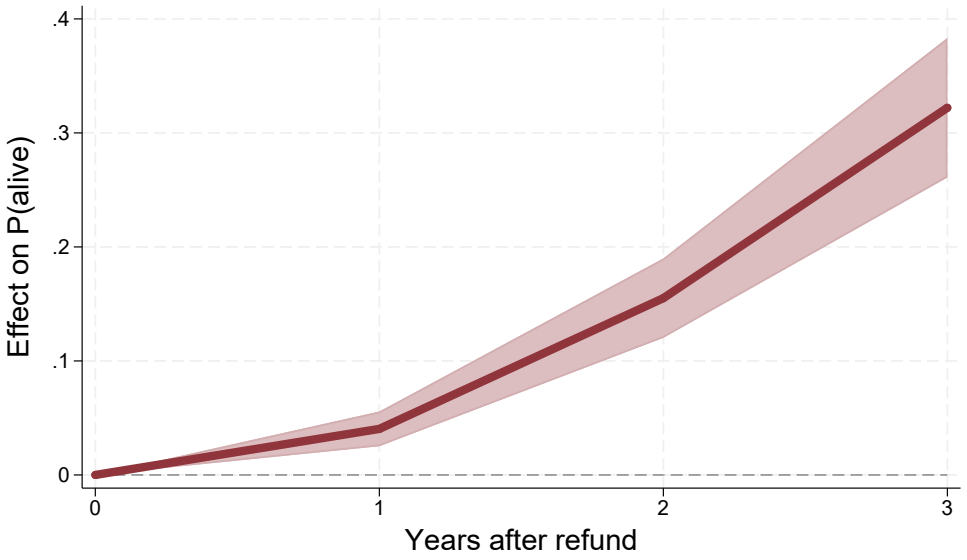
where $\kappa_{c(i),h}$ are fixed effects for the exact matching cells of Section 6.3 and X_i collects the baseline covariates used in the match. The coefficient γ_h absorbs the relationship between survival and the entitlement that is common to corporations and pass-throughs, which is to say the selection of firms with strong profit histories into large entitlements, whatever its source. The coefficient of interest is β_h, the differential gradient for corporations, which under the identifying assumption of Section 6.1 is the causal effect of the carryback on the probability of surviving h years, per dollar of entitlement per dollar of payroll. Standard errors are clustered by matched cell, the level at which treatment is effectively assigned, since every corporation and its matched pass-through share an industry and size class and are therefore exposed to the same local shocks. Because D_i is scaled by payroll, magnitudes are obtained by evaluating β_h at the average entitlement among corporations, roughly 5 percent of baseline payroll.

7.2 Main results

Figure 5 reports the estimates of β_h at horizons of one to three years. The differential gradient is positive at every horizon and grows with distance from the refund, rising from roughly 0.04 one year out to 0.15 at two years and 0.32 at three. At the average entitlement, the horizon-three estimate implies a 1.6 percentage point increase in the probability of surviving three years, relative to a baseline three-year survival rate of 85.2 percent. We also estimate equation (12) with the outcome

equal to one only if the firm is alive and reports positive taxable income, the margin that determines whether a saved firm resumes remitting tax; Appendix Figure B.2 reports it: the effect on operating with positive taxable income is likewise positive and grows with the horizon.

Figure 5: The effect of carrybacks on firm survival



Notes: The figure plots the coefficients β_h from equation (12) at horizons of one to three years after the expected refund, with 95 percent confidence intervals, estimated on the matched panel of corporations and pass-throughs. Standard errors are clustered by matched cell. Statistics are rounded in accordance with ABS disclosure rules.

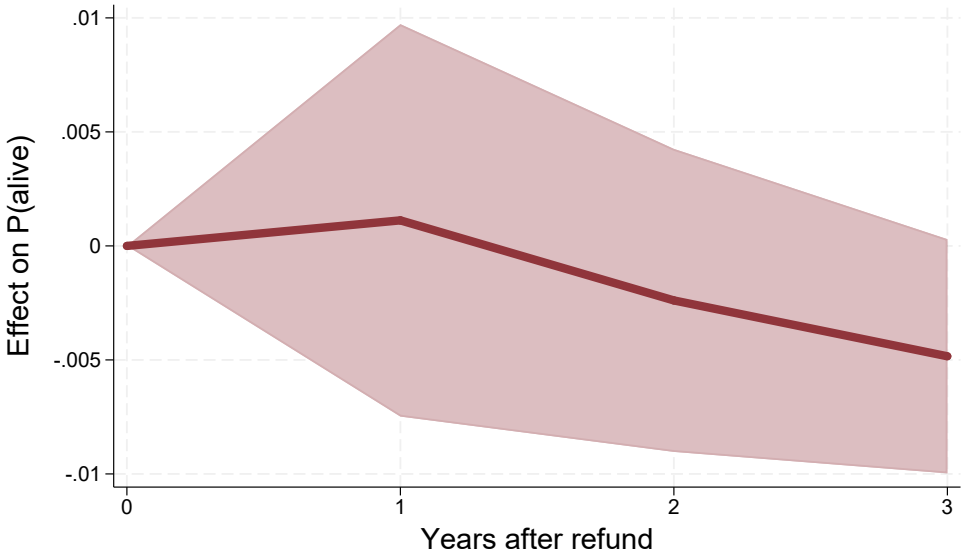
7.3 A placebo policy in 2014-15

The identifying assumption would fail if corporations with large entitlements fared differently from pass-throughs with the same entitlements for reasons unrelated to the refund. The natural test is to rerun the entire construction in a year in which the same legal forms and the same kinds of profit histories were present but no carryback existed. We use 2014-15, a year comfortably inside our panel, after the repeal of the 2012-13 provision described in Section 3 and well before the pandemic policy was announced. We apply the statutory formula to 2014-15 losses to construct a synthetic entitlement for every firm, corporations and pass-throughs alike, rebuild the matched sample on 2013-14 baseline characteristics, and re-estimate equation (12) at the same horizons. No refunds were paid against these entitlements, so under the identifying assumption the differential gradient should be zero; a nonzero β_h would indicate that legal form and profit history interact independently of the policy. Figure 6 reports the estimates. The differential gradient is a tight zero at every horizon: the point estimates are smaller than 0.005 in absolute value, two orders of magnitude below the estimates of Figure 5, and never statistically distinguishable from zero.

7.4 Discussion

Two features of the estimates are worth drawing out. First, the magnitude is consistent with the liquidity mechanism of Section 2. A 1.6 percentage point increase against a baseline three-year exit rate of roughly 15 percent implies that about one in nine exits among firms receiving the average entitlement was averted: the refund was pivotal for a minority of recipients, as the framework

Figure 6: Placebo: synthetic carrybacks in 2014-15



Notes: The figure plots the coefficients β_h from equation (12) estimated on synthetic entitlements constructed by applying the statutory formula to 2014–15 losses, a year in which no carryback was available, with the matched sample rebuilt on 2013–14 baseline characteristics. Shaded bands are 95 percent confidence intervals. Statistics are rounded in accordance with ABS disclosure rules.

predicts when only the firms near their cash floor are moved across it. Second, the effect grows with the horizon rather than fading, which is what saving a firm, rather than delaying its exit, looks like: a refund that merely postponed closure would generate a gradient that peaks early and declines. The placebo of Figure 6 argues that the gradient is not a general feature of the interaction of legal form with profit history. The survival margin also frames the employment responses of Section 8: because exiting firms remain in the monthly panel at zero payroll and employment, part of the response there operates through firms that continue to exist at all.

8 The effects of carrybacks on employment and payroll

Survival is an annual outcome, measured from records that resolve at the financial year. The remaining outcomes require sharper timing: a refund arrives in a particular month, and payroll can adjust within weeks of it. From this point the analysis therefore moves from annual to monthly data. We set aside the annual tax records that have organized the paper so far and use the monthly payroll reporting described in Section 4, which records the number of employees each firm pays, the gross amounts paid to them, and the tax withheld, in every month from 2019-20 onward. The monthly frequency is what allows outcomes to be aligned to the arrival of the refund itself, and it is the basis of everything in this section.

8.1 Timing and specification

A carryback was monetized when the firm lodged the return containing the claim, so the refund's arrival is governed by the lodgement calendar rather than by the loss. We define the expected refund month m^* as nine months after the close of the loss year, consistent with typical lodgement timing,

and measure event time $k = m - m^*$ in months. For outcome $Y_{i,m}$ of firm i in calendar month m, and for $k \in \{-16, \ldots, 16\} \setminus \{0\}$, we estimate

$$Y_{i,m} = \alpha_i + \lambda_m + \sum_k \gamma_k \, CB_i \cdot \mathbf{1}\{m - m^* = k\} + \sum_k \beta_k \, \text{Corp}_i \times CB_i \cdot \mathbf{1}\{m - m^* = k\} + \varepsilon_{i,m}, \quad (13)$$

where α_i are firm fixed effects and λ_m calendar month fixed effects. The treatment CB_i is the potential carryback in dollars, unscaled, so that when the outcome is also in dollars β_k reads directly as the effect of a dollar of carryback. The expected refund month $k = 0$ is omitted, so every coefficient is measured relative to it. As in the survival design, γ_k absorbs the relationship between the entitlement and the outcome path common to both legal forms, and the differential path β_k traces the causal effect of the refund month by month under the assumption of Section 6. The pre-period coefficients are the design's internal test: for $k < 0$ the refund has not arrived, and β_k should be flat at zero if the counterfactual gradients are parallel. The sample is the matched panel; firms that exit remain in the panel at zero payroll and zero employment, so the estimates embed the survival margin rather than conditioning on it. Standard errors are clustered by firm. The outcomes are total monthly payroll in dollars, employment, and average payment per employee; we re-estimate in logs and with entropy reweighting in place of the matched pairing as robustness checks.

8.2 Payroll and employment

Figure 7 reports the differential path β_k for total monthly payroll, and Figure 8 for employment. Both differential paths are flat in the sixteen months before the refund and turn up after it. A dollar of carryback is associated with roughly ten cents of additional monthly payroll in the year following the refund, and at the average refund the employment response builds to roughly 0.8 additional employees per firm. Neither response fades by the end of the sixteen-month window. The responses are similar in logs and when the matched pairing is replaced by entropy reweighting; Appendix Figures B.3 through B.6 report these checks.

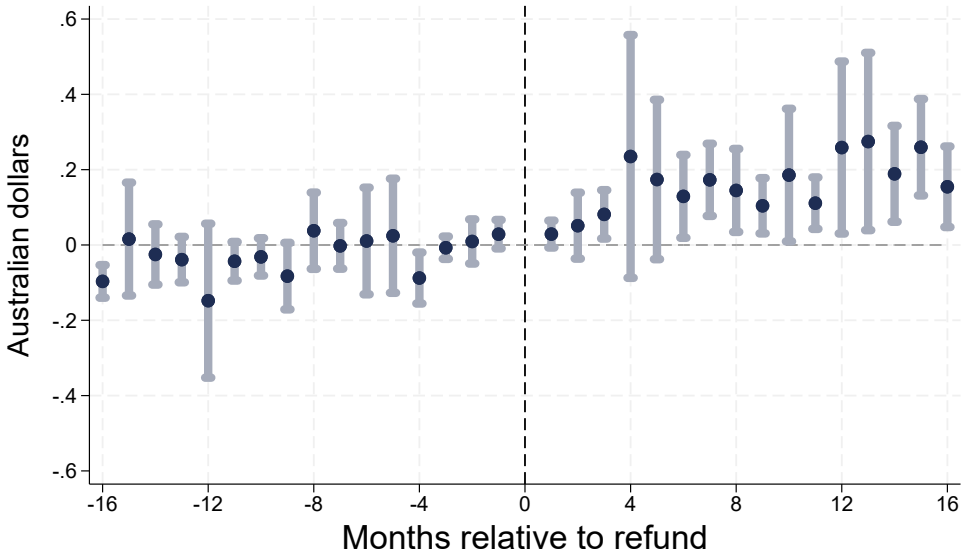
8.3 Pay per employee

The payroll response could reflect more workers, better-paid workers, or both, and the split matters for incidence: higher pay accrues to workers who would have been employed in any case, while an employment response operates through jobs that would otherwise not exist. Figure 9 reports equation (13) with average payment per employee as the outcome. The differential path is a tight zero throughout, before and after the refund: the refund changed how many workers treated firms paid, not how much they paid each of them. The employment margin, not the wage margin, carries the payroll response.

8.4 Retention versus new hiring

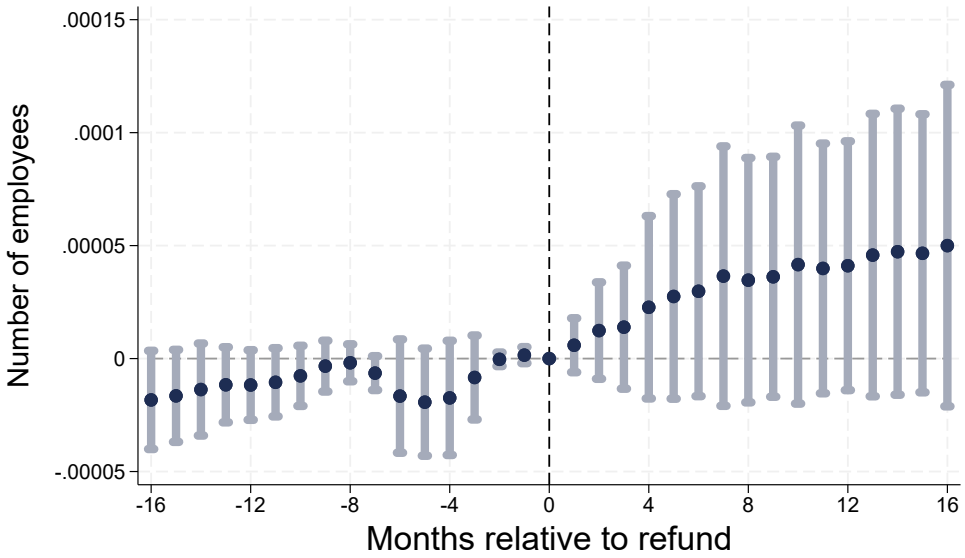
An employment response can itself operate on two margins: the retention of incumbent workers who would otherwise have been let go, or the hiring of new ones. The distinction carries the interpretation of the policy. Retention preserves existing firm-worker matches and shields incumbent workers from displacement, the insurance reading suggested by the framework of Section 2; hiring is expansion, closer to an investment or stimulus reading. We therefore split each firm's monthly payroll by worker tenure, $Y_{i,m} = Y^{\text{ret}}_{i,m} + Y^{\text{new}}_{i,m}$, classifying a payment as going to a new worker if that worker is absent from the firm's payroll in every earlier month and to a retained worker otherwise, and

Figure 7: The effect of carrybacks on monthly payroll



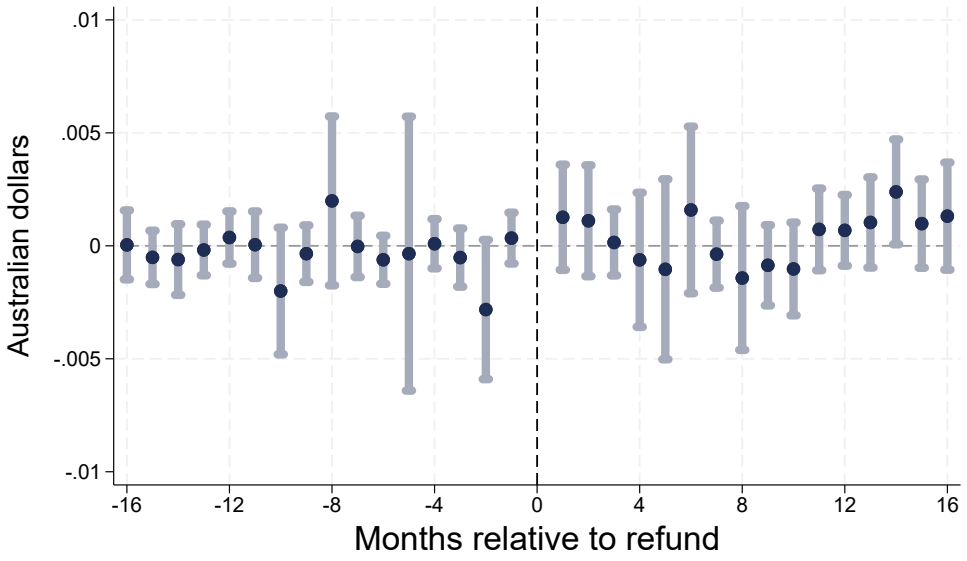
Notes: The figure plots the differential path β_k from equation (13) with total monthly payroll in dollars as the outcome, for event months $k \in \{-16, \ldots, 16\}$ relative to the expected refund, with 95 percent confidence intervals.

Figure 8: The effect of carrybacks on employment



Notes: The figure plots the differential path β_k from equation (13) with the number of employees paid as the outcome, for event months $k \in \{-16, \ldots, 16\}$ relative to the expected refund, with 95 percent confidence intervals.

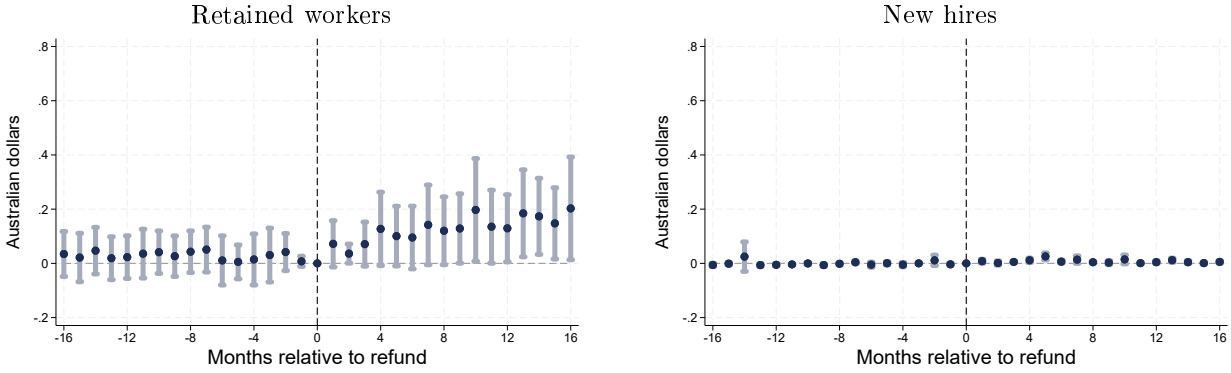
Figure 9: The effect of carrybacks on average pay per employee



Notes: The figure plots the differential path β_k from equation (13) with average payment per employee in dollars as the outcome, for event months $k \in \{-16, \ldots, 16\}$ relative to the expected refund, with 95 percent confidence intervals.

estimate equation (13) separately for the two components. Figure 10 reports both on a common vertical scale. The payroll response is concentrated among retained workers: the retained-worker path mirrors the total payroll response of Figure 7, while the new-hire path is flat and precisely estimated throughout; Appendix Figure B.7 shows it on its own scale. The refund preserved existing employment relationships rather than funding expansion.

Figure 10: Decomposition of the payroll response: retained workers and new hires



Notes: The figure plots the differential path β_k from equation (13) estimated separately for payroll paid to retained workers (left panel) and to new hires (right panel), on a common vertical scale, with 95 percent confidence intervals. Statistics are rounded in accordance with ABS disclosure rules.

9 Conclusion

Australia's introduction of tax loss carrybacks during the COVID-19 pandemic is a natural experiment in narrowing the asymmetry at the heart of the corporate income tax. Corporations could convert a current loss into an immediate refund of tax paid on past profits; the pass-through businesses operating alongside them could not. Using the universe of business tax returns linked to monthly payroll records, we reconstruct every firm's statutory entitlement and compare corporations to pass-throughs assigned the counterfactual refund their own histories would have commanded.

Three findings emerge. First, the eligibility criterion is a powerful tag: conditioning on prior tax paid directs refunds toward firms with long histories of profit, which return to profitability at roughly twice the rate of the average loss-maker, and take-up, conditional on the size of the entitlement, is highest among financially constrained firms. Second, the refund had real effects: it raised three-year survival by roughly 1.6 percentage points at the average entitlement and produced persistently higher employment and payroll, with no detectable change in pay per employee. Third, the payroll response operates through the retention of incumbent workers rather than through new hiring. Taken together, the estimates describe a policy that functioned as insurance: it reached viable firms hit by a transitory shock, was claimed where cash was scarcest, and preserved employment relationships that would otherwise have dissolved, achieving objectives similar to crisis policies that target jobs directly while operating automatically through the existing tax return.

Two questions remain open. The framework of Section 2 implies that the fiscal cost of the policy depends on the rate at which refunded losses would have been deducted in the counterfactual, and that the same estimates, combined with the earnings of the workers retained, aggregate into a welfare evaluation of the reform. And the firm-level payroll responses documented here bound, but do not identify, the earnings workers would have lost absent the refund, which requires following workers across employers. We take up both in ongoing work.

References

Rosanne Altshuler and Alan J. Auerbach. The significance of tax law asymmetries: An empirical investigation. *Quarterly Journal of Economics*, 105(1):61–86, 1990. 1

Rosanne Altshuler, Alan J. Auerbach, Michael Cooper, and Matthew Knittel. Understanding U.S. corporate tax losses. In Jeffrey R. Brown and James M. Poterba, editors, *Tax Policy and the Economy, Volume 23*, pages 73–122. University of Chicago Press, Chicago, 2009. 1

Alan J. Auerbach. The dynamic effects of tax law asymmetries. *Review of Economic Studies*, 53 (2):205–225, 1986. 1

Alan J. Auerbach and James M. Poterba. Tax loss carryforwards and corporate tax incentives. In Martin Feldstein, editor, *The Effects of Taxation on Capital Accumulation*, pages 305–342. University of Chicago Press, Chicago, 1987. 1

David Autor, David Cho, Leland D. Crane, Mita Goldar, Byron Lutz, Joshua Montes, William B. Peterman, David Ratner, Daniel Villar, and Ahu Yildirmaz. An evaluation of the paycheck protection program using administrative payroll microdata. *Journal of Public Economics*, 211: 104664, 2022. 1

Inga Bethmann, Martin Jacob, and Maximilian A. Müller. Tax loss carrybacks: Investment stimulus versus misallocation. *The Accounting Review*, 93(4):101–125, 2018. 1

Ricardo J. Caballero, Takeo Hoshi, and Anil K. Kashyap. Zombie lending and depressed restructuring in Japan. *American Economic Review*, 98(5):1943–1977, 2008. 1

Pierre Cahuc, Francis Kramarz, and Sandra Nevoux. The heterogeneous impact of short-time work: From saved jobs to windfall effects. IZA Discussion Paper 14381, Institute of Labor Economics (IZA), 2021. 1

Michael Cooper and Matthew Knittel. Partial loss refundability: How are corporate tax losses used? *National Tax Journal*, 59(3):651–663, 2006. 1

Christine L. Dobridge. Tax loss carrybacks as firm fiscal stimulus: A tale of two recessions. *The Accounting Review*, 96(6):153–181, 2021. 1

Daniel Dreßler and Michael Overesch. Investment impact of tax loss treatment—empirical insights from a panel of multinationals. *International Tax and Public Finance*, 20(3):513–543, 2013. 1

Jesse Edgerton. Investment incentives and corporate tax asymmetries. *Journal of Public Economics*, 94(11–12):936–952, 2010. 1

Jack Favilukis, Ronald Giammarino, Jose Pizarro, and Mikhail Simutin. Tax-loss carry forwards and firms' risk. Working paper, available at SSRN 2812410, 2025. 2.1, 2.1

Giulia Giupponi and Camille Landais. Subsidizing labour hoarding in recessions: The employment and welfare effects of short-time work. *Review of Economic Studies*, 90(4):1963–2005, 2023. 1

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

Pierre-Olivier Gourinchas, Şebnem Kalemli-Özcan, Veronika Penciakova, and Nick Sander. SME failures under large liquidity shocks: An application to the COVID-19 crisis. *Journal of the European Economic Association*, 23(2):673–718, 2025. 1, 2.1

João Granja, Christos Makridis, Constantine Yannelis, and Eric Zwick. Did the paycheck protection program hit the target? *Journal of Financial Economics*, 145(3):725–761, 2022. 1

Richard C. Green and Eli Talmor. The structure and incentive effects of corporate tax liabilities. *Journal of Finance*, 40(4):1095–1114, 1985. 2.1

Lisa Hillmann and Martin Jacob. The effect of limited tax loss carryforwards on corporate investment. *Journal of Accounting and Economics*, 80(1):101724, 2025. 1

Patrick J. Kennedy, Christine Dobridge, Paul Landefeld, and Jacob Mortenson. The efficiency-equity tradeoff of the corporate income tax: Evidence from the Tax Cuts and Jobs Act. Working paper, 2026. 1, 6.1

Reinald Koch, Svea Holtmann, and Henning Giese. Losses never sleep—the effect of tax loss offset on stock market returns during economic crises. *Journal of Business Economics*, 93(1):59–109, 2023. 1

Dominika Langenmayr and Rebecca Lester. Taxation and corporate risk-taking. *The Accounting Review*, 93(3):237–266, 2018. 1

Alexander Ljungqvist, Liandong Zhang, and Luo Zuo. Sharing risk with the government: How taxes affect corporate risk taking. *Journal of Accounting Research*, 55(3):669–707, 2017. 1, 2.2

Martin C. Schmalz, David A. Sraer, and David Thesmar. Housing collateral and entrepreneurship. *Journal of Finance*, 72(1):99–132, 2017. 2.1

Danny Yagan. Capital tax reform and the real economy: The effects of the 2003 dividend tax cut. *American Economic Review*, 105(12):3531–3563, 2015. 1, 6.1

Eric Zwick. The costs of corporate tax complexity. *American Economic Journal: Economic Policy*, 13(2):467–500, 2021. 1, 2.1, 5.3

Eric Zwick and James Mahon. Tax policy and heterogeneous investment behavior. *American Economic Review*, 107(1):217–248, 2017. 1

Appendix

A Data Appendix

This appendix documents the construction of the analysis data set. References to numbered BLADE tables follow the Australian Bureau of Statistics*BLADE Modular Product Data Item List, 2001-02 to 2024-25*(released September 2025), and variable names are reproduced exactly as they appear in that document. Item and label references on the tax return are to the 2020-21 company tax return unless otherwise noted.

A.1 Source modules

Table A.1 lists the BLADE modules we draw on, the level at which each is reported, its frequency and coverage, and the items we take from it. Every module is delivered as one file per financial year, indexed by the two-digit time series identifier `tsid`, which records the last two digits of the end year of the financial year, so that the 2020-21 year carries the value 21. Quarterly and monthly files additionally carry `quarter` and `month_actioned`; we collapse these to the financial year for all variables except the payroll series used in the monthly event studies, which we retain at monthly frequency.

A.2 Identifiers, linkage and panel structure

BLADE carries two unit identifiers. The de-identified Australian Business Number, `bn`, identifies the tax-filing entity; all of the administrative modules (Tables 4, 5, 6, 7 and 26) are reported at this level. The de-identified economic unit identifier, `id`, identifies the enterprise, and is the unit at which the Bureau constructs its business statistics. A third identifier, `bg_id`, identifies the BLADE enterprise group; entities with a non-missing value belong to the Bureau's profiled population, for which unit structure is actively maintained rather than inferred from registration data.

We link `bn` to `id` using the cross-sectional ID to BN Key (Appendix 1 of the data item list), merging on `bn` and `tsid`. The key carries two flags describing the mapping. We drop observations with `one_abn_to_many_tau` $= 1$, that is, filing entities whose activity the Bureau splits across more than one economic unit: the tax return and payroll records for such an entity are reported as a single undivided total and cannot be apportioned. We retain observations with `many_abn_to_one_tau` $= 1$; these are precisely the multi-entity enterprises that motivate the aggregation step in Section A.4. The key also carries a `match` variable describing how each entity was assigned to a unit, ranging from an exact assignment where the enterprise group contains a single unit through to matches on progressively coarser industry classifications. We do not condition on it, but note that the coarser match categories are concentrated among small entities within large groups.

The resulting panel is unbalanced. Entities enter when they first register and lodge, and leave when they cease to lodge. We do not impose balance, and we treat the last month in which a firm reports positive payroll as the basis for the survival measures described in the main text.

A.3 Sector and legal-form exclusions

From the cross-sectional Indicative module (Table 1) we drop non-profit institutions, identified by `x_npi` $= 1$, and public-sector entities, identified by `x_sector` $= 2$. We additionally drop entities whose type of legal organization, `x_tolo`, falls in the government ranges: codes 21 to 26 (Commonwealth departments, legislatures and courts, territory municipal authorities, marketing boards,

Table A.1: BLADE modules used in this paper

Module	BLADE table	Reporting unit	Frequency (coverage)	Content used here
Cross-sectional Indicative	Table 1	`bn`, `id`	Quarterly (2001–02 to 2024–25)	Industry (`d_div06`, `x_anzsic06`), legal form (`x_tolo`), institutional sector (`d_sisca08`), state (`x_state`), non-profit and public-sector flags (`x_npi`, `x_sector`), activity status (`x_al_st`), enterprise group (`bg_id`)
Business Activity Statement	Table 4	`bn`	Quarterly and annual (2001–02 to 2024–25)	`turnover`, `capex`, `oexp`, `tot_expenses`, `wages`, `payg_tax_withheld`, `gst_payable`, `exports_amt`
Pay As You Go	Table 5	`bn`	Annual (2001–02 to 2022–23)	Headcount (`hcnt`) and full-time equivalent employment (`fte`)
Business Income Tax	Table 6	`bn`	Annual (2001–02 to 2022–23)	Company, individual, partnership and trust return items; see Table A.3
Single Touch Payroll	Table 7	`bn`	Monthly (2019–20 to 2024–25)	Employees paid (`d_total_payees`), gross payments (`ed_incm_grs`, `ed_pmt_sumry_totl_grs_pmt`), withholding (`ed_tax_wheld_totl`), reference month (`month_actioned`)
Business Birthdate	Table 26	`bn`	Static (1992–93 to 2023–24)	`birth_date`, used to construct firm age
ID to BN Key	Appendix 1	`bn`, `id`	Annual (2001–02 to 2024–25)	`id`, `one_abn_to_many_tau`, `many_abn_to_one_tau`, `match`

Notes: Table numbers refer to the ABS *BLADE Modular Product Data Item List, 2001–02 to 2024–25* (released September 2025). `bn` is a de-identified Australian Business Number; `id` is a de-identified ABS economic unit identifier. We do not use the ABS survey modules (Tables 8–22), the business location, insolvency and intellectual property modules (Tables 23–27 other than Table 26), or the requestable modules (Tables 28–48).

statutory authorities, and other Commonwealth entities), 31 to 36 (state departments, legislatures and courts, marketing boards, local government authorities, other local government, and other state entities), and 41 to 42 (diplomatic and trade missions, other foreign governments). As a cross-check we drop entities whose derived institutional sector, `d_sisca08`, takes the values 3000 (general government), 5000 (non-profit institutions serving households) or 6000 (rest of the world). The three criteria overlap heavily; applying all of them rather than any one is a conservative choice that costs few genuine operating businesses.

We do not exclude entities on the basis of legal form beyond this. The comparison between corporations and pass-throughs requires that trusts (code 9), family and other partnerships (codes 7 and 8) and sole proprietors (code 6) be retained, and the corporate population comprises codes 1 through 5, that is, proprietary, limited, proprietary limited, no liability, and other registered companies. In the non-profiled population the Bureau collapses several of these codes: code 1 is folded into code 3, code 4 into code 2, and code 11 into code 5. We therefore work with the collapsed classification throughout, so that the corporate indicator has a consistent meaning across the profiled and non-profiled populations.

Legal form as recorded on the business register is imperfect, and we cross-check it against the return form each entity actually lodged. BLADE Table 6 carries a separate flag for each of the four return types (`bit_comp_yyyy`, `bit_ind_yyyy`, `bit_part_yyyy`, `bit_trust_yyyy`). We classify an entity as a corporation in a given year if it lodged a company return that year, and fall back on the register's legal form only where no return is observed.

A.4 Aggregation from the filing entity to the firm

Having formed a panel indexed by `bn` $\times$ `tsid` and merged in the tax, activity statement, payroll and birth date modules, we collapse to `id` $\times$ `tsid`. Table A.2 gives the rule applied to each class of variable. The great majority of economic units contain a single filing entity, so these rules bind only for the multi-entity enterprises identified in the previous subsection; we have verified that the results in the paper are unchanged if the analysis is run at the filing-entity level on the subsample of single-entity firms.

A.5 Variable dictionary

Table A.3 maps each concept used in the paper to the corresponding item on the company income tax return, organized by the panel of the return on which it appears. Label positions shift somewhat across years, but the BLADE variable names are stable, and where a variable was renamed the data item list records both names. Table A.4 gives the corresponding items on the individual, partnership and trust forms. Table A.1 above lists the items taken from the non-tax modules.

Table A.2: Aggregation rules from the filing entity to the firm

Rule	Variables	Rationale
Sum across constituents	All dollar flows and stocks from Tables 4 and 6; employment counts `hcnt`, `fte`, `d_total_payees`; carryback amounts and offsets	The enterprise's income, expenses, balance sheet and workforce are the sum of those of its filing entities
Largest constituent	`d_div06`, `x_anzsic06`, `c_anzsic`, `x_tolo`, `x_state`, `x_pcode`	Categorical attributes cannot be summed; we take them from the constituent with the largest total income, falling back on Business Activity Statement turnover where total income is missing
Earliest value	`birth_date`	Firm age is the age of the oldest constituent, which is when the enterprise began operating
Any constituent	`c_bm_sgmt_cd`, `c_base_rt_bus_cd`, `c_sobconin`, form indicators `bit_comp_yyyy` and its analogues, an indicator for having claimed a carryback	Status indicators are set to one if any constituent reports them
Recomputed after aggregation	Statutory tax rate, potential carryback, profit margins, revenue and payroll per worker, leverage, cash-to-assets	Ratios of aggregates rather than aggregates of ratios

Notes: Aggregation is performed within `id` $\times$ `tsid` cells. The vast majority of economic units contain a single filing entity, so these rules bind only for the multi-establishment enterprises identified by `many_abn_to_one_tau` $= 1$.

Table A.3: Company income tax return items used from BLADE Table 6

Concept in the paper	BLADE variable	Return label and notes
Panel A. Identifiers and firm characteristics		
Filing entity	`bn`	De-identified Australian Business Number
Financial year	`tsid`	Last two digits of the end year; 2020–21 = 21
Company form indicator	`bit_comp_yyyy`	Set to 1 where a company return exists; analogous flags `bit_ind_yyyy`, `bit_part_yyyy`, `bit_trust_yyyy` identify the three pass-through forms
Industry	`c_anzsic`	ANZSIC code reported on the return
Small business entity	`c_bm_sgmt_cd`	Item 3, label F1; code 5 = yes
Base rate entity	`c_base_rt_bus_cd`	Item 3, label F2; code 5 = yes. Determines the reduced rate from 2017–18
Consolidation status	`c_sobconin`	Item 3, labels Z1–Z2; head company or subsidiary member of a consolidated group
Significant global entity	`c_clnt_glbl_ent_typ_cd`	Item 3, label G1
Aggregated turnover	`c_agrgtd_tnovr_amt`	Item 8, label Y. Used for the $5 billion carryback eligibility test and the base rate entity threshold
Panel B. Income statement (item 6)		
Sales	`c_sales`	Other sales of goods and services, label C
Interest income	`c_grosintr`	Gross interest, label F
Rent and leasing income	`c_grosrent`	Gross rent and other leasing and hiring income, label G
Dividends received	`c_grosdivd`	Total dividends, label H
Government industry payments	`c_govind`	Assessable government industry payments, label Q. Captures JobKeeper and cash-flow boost receipts
Total income	`c_totlinc`	Label S
Cost of sales	`c_costsale`	Label A (expenses)
Contractor and commission expenses	`c_extnlabr`	Label C (expenses)
Superannuation expenses	`c_superann`	Label D (expenses)
Rent and lease expenses	`c_rentexps`, `c_leaexpau`, `c_leaexpos`	Labels H, F and I (expenses)
Interest expenses	`c_intexpau`, `c_intexpos`	Labels V and J (expenses)
Depreciation expenses	`c_deprexps`	Label X (expenses)
Total expenses	`c_totlexps`	Label Q (expenses)
Accounting profit	`c_toprolos`	Total profit or loss, label T
Operating profit	`c_oproloss`	Operating profit or loss
Panel C. Balance sheet (item 8)		
Opening and closing stock	`c_opnstock`, `c_clostock`	Labels A and B

continued on next page

Table A.3, continued

Concept in the paper	BLADE variable	Return label and notes
Trade debtors and creditors	`c_debtors`, `c_creditrs`	Labels C and F
Current and total assets	`c_currasst`, `c_totlasst`	Labels D and E
Current and total liabilities	`c_currliab`, `c_totlliab`	Labels G and H
Total debt	`c_totl_debt_amt`	Label J
Shareholders' funds	`c_sharhold`	
Total wage bill	`c_totlwage`	Total salary and wage expenses, label D (expenses)
Cash proxy	`c_currasst` − `c_debtors` − `c_clostock`	Cash is not reported; we net inventories and receivables out of current assets
Panel D. Reconciliation to taxable income (item 7)		
Net capital gain	`c_ncg_amt`	Label A
Exempt income	`c_exemptax`	Label V (expenses)
Non-deductible exempt income expenditure	`c_nondedex`	Label U. Netted against `c_exemptax` to form net exempt income
Non-deductible expenses	`c_non_ddctbl_expnss_amt`	Label W
Tax losses deducted	`c_ddctd_tl_amt`	Label R (expenses)
Tax losses transferred in	`c_forgn_fincl_ent_tl_tfrd_in_amt`	Label S (expenses)
Taxable income or loss	`c_taxinc`	Label T (expenses). Negative in loss years
Tax losses carried forward	`c_tlcf_amt`	Item 13, label U
Net capital losses carried forward	`c_nlcf_amt`	Item 13, label V
Panel E. Calculation statement		
Taxable or net income	`c_taxicalc`	Label A
Tax on taxable income	`c_taxti`	Label T1. Numerator of the implied statutory rate
Gross tax	`c_grstx`	Label B
Non-refundable offsets	`c_nonrfnblencfto`, `c_nonrfnblefcfto`	Labels C and D
Refundable tax offsets	`c_rfnbltaxo`	Label E. Contains the loss carryback offset
Tax payable	`c_subt5`	Label T5
PAYG instalments raised	`c_pygis_rsd_amt`	Label K
Amount due or refundable	`c_tax_pybl_or_rfndbl_totl_amt`	Label S
Panel F. Imputation and distributions (item 8)		
Franked dividends paid	`c_divfrank`	Label J
Unfranked dividends paid	`c_divufran`	Label K
Franking credits received	`c_divs_fcr_amt`	Item 7, label J
Opening franking account balance	`c_div_frnkg_acnt_opng_bal_amt`	Label P
Closing franking account balance	`c_div_frnkg_acnt_closg_bal_amt`	Caps the carryback refund; `c_franbalc` (label M) is used where the closing balance is missing

continued on next page

Table A.3, continued

Concept in the paper	BLADE variable	Return label and notes
Excess franking offsets	`c_divs_frnkg_tax_ofsts_excs_amt`	Label H (expenses)
Panel G. Loss carryback items (2020–21 to 2022–23 returns, item 13)		
Loss carried back, by pairing	`c_fy_`Y_L`_to_fy_`Y_G`_tax_lss_cba`	One variable for each of the ten admissible loss-year Y_L to gain-year Y_G pairings; see Table A.6
Corporate tax rate in the loss year	`c_fy_2020_tax_rt`, `c_fy_2021_tax_rt`, `c_fy_2022_tax_rt`	Reported to four decimal places
Net exempt income of the gain year	`c_fy_`Y_G`_net_exmt_incm_amt`	Available for $Y_G \in \{2018\text{–}19, \ldots, 2021\text{–}22\}$
Income tax liability of the gain year	`c_fy_`Y_G`_it_lblty_amt`	Caps the offset attributable to that gain year
Loss carryback tax offset	`c_lss_crry_bck_to_amt`, `c_lss_crry_bck_to`	Item 13, label S. The refund actually claimed
Early balancer items	`c_chsnyrtaxlssbckerlyyramt`, `c_chsnmdyrntutlstxlsbckerlyyramt`, `c_erlyst_yr_net_exmt_incm_amt`, `c_erlyst_yr_it_lblty_amt`	Used by substituted-accounting-period filers in place of the year-specific labels
Panel H. Investment incentives (items 9 and 10)		
Temporary full expensing deduction	`c_tfe_bus_ddctn_amt`	Item 9, label S
Instant asset write-off	`c_wrt_off_dtl_itm_amt`	Item 9, label O
Backing business investment deductions	`c_frst_yr_ast_dprtn_bbi_ddctn_to`, `c_sbsq_yr_ast_dprtn_bbi_ddctn_to`	Item 9, labels M and N
Depreciating assets first deducted	`c_intdepas`, `c_othdepas`	Item 9, labels A and B

Table A.4: Corresponding items on the pass-through return forms

Concept	Individual form	Partnership form	Trust form
Form indicator	bit_ind_yyyy	bit_part_yyyy	bit_trust_yyyy
Industry	i_anzsic	p_anzsic	—
Cost of sales	i_costsale	p_costsale	—
Total expenses	i_totlexps	p_totlexps	—
Total wage bill	i_totlwage	p_totlwage	—
Depreciation	i_deprexps	p_deprexps	—
Trade debtors	i_debtors	p_debtors	—
Total assets	—	p_totlasst	—
Total liabilities	—	p_totlliab	—
Net income or loss	i_netincls, i_ntinnpp	p_netilbus	t_netilbus
Net Australian income or loss	—	p_ausn_incm_or_lss_net_amt	t_ausn_incm_or_lss_net_amt
Total net income or loss	—	p_incm_or_lss_net_totl_amt	t_incm_or_lss_net_totl_amt
Aggregated turnover	—	p_agrgtd_tnovr_amt	t_agrgtd_tnovr_amt
Tax losses deducted	—	—	t_ddctd_tl_amt
Tax losses carried forward	—	—	t_tlcf_amt
Temporary full expensing	i_tfe_bus_ddctn_amt	p_bus_ddctn_amt	—

Notes: The pass-through forms are considerably less detailed than the company form; in particular none of the three reports a calculation statement, a franking account, or any loss carryback item, since pass-through entities were statutorily ineligible for the carryback. Where a balance sheet item is unavailable on the individual or trust form we fall back on the Business Activity Statement equivalents from BLADE Table 4. A dash indicates that no corresponding item exists on that form.

A.6 Derived variable: the statutory tax rate

Let τ_{it} denote the statutory corporate tax rate facing firm i in year t. Our preferred measure is the rate implied by the return itself, the ratio of tax charged on taxable income to taxable income:

$$\hat{\tau}_{it} \;=\; \frac{\texttt{c_taxti}_{it}}{\texttt{c_taxinc}_{it}}.$$

We accept $\hat{\tau}_{it}$ as the statutory rate when taxable income is strictly positive and $\hat{\tau}_{it}$ lies within 0.001 of one of the two rates in force in year t, as listed in Table A.5.

The ratio is unavailable or uninformative in three cases. It is undefined when taxable income is zero or negative, which is by construction true in exactly the loss years that matter most for the carryback calculation. It is contaminated when the firm is taxed under one of the special schedules -- life insurance companies, retirement savings account providers, pooled development funds, credit unions, and not-for-profit companies each face a different structure, and in the not-for-profit and credit union cases the schedule is not even proportional. And it is occasionally implausible for reasons we cannot diagnose, presumably reflecting amendments and reporting error. In each case we impute the rate from Table A.5 using the firm's own reported status: `c_base_rt_bus_cd` (code 5 indicates a base rate entity) for 2017-18 onward, and `c_bm_sgmt_cd` (code 5 indicates a small business entity) for 2015-16 and 2016-17, in each case checking the reported status against aggregated turnover, `c_agrgtd_tnovr_amt`, and the threshold in force that year. For years through 2014-15 the rate was a flat 30 percent and no imputation is required. Where the status indicator is itself missing we assign the reduced rate if aggregated turnover is below the threshold and the standard rate otherwise.

One subtlety deserves note. From 2017-18 the reduced rate applies to base rate entities, which are defined not only by a turnover threshold but also by the requirement that no more than 80 percent of assessable income be passive -- dividends and the franking credits attached to them, royalties, rent, interest, gains on qualifying securities, and net capital gains. Turnover alone therefore does

not determine the rate. This is why we prefer the implied rate wherever it is available, and why the imputation leans on the firm's own reported base rate entity status rather than reconstructing the passive income test from the income items on the return.

Table A.5: Statutory company tax rates used for imputation

Income year	Reduced-rate concept	Turnover threshold	Rate: reduced / standard
2001–02 to 2014–15	—	—	30% / 30%
2015–16	Small business entity	\$2m	28.5% / 30%
2016–17	Small business entity	\$10m	27.5% / 30%
2017–18	Base rate entity	\$25m	27.5% / 30%
2018–19	Base rate entity	\$50m	27.5% / 30%
2019–20	Base rate entity	\$50m	27.5% / 30%
2020–21	Base rate entity	\$50m	26.0% / 30%
2021–22	Base rate entity	\$50m	25.0% / 30%
2022–23	Base rate entity	\$50m	25.0% / 30%

Notes: Rates are published by the Australian Taxation Office. From 2017–18 the reduced rate applies to base rate entities, defined as companies with aggregated turnover below the threshold and with no more than 80% of assessable income classified as base rate entity passive income; for 2015–16 and 2016–17 it applied to small business entities. Special rates apply to life insurance companies, retirement savings account providers, pooled development funds, credit unions and not-for-profit companies; firms whose implied rate matches one of these schedules are flagged and, in the case of not-for-profit companies, dropped at the sample selection stage. We use `c_base_rt_bus_cd` (code 5 = base rate entity) from 2017–18 and `c_bm_sgmt_cd` (code 5 = small business entity) for 2015–16 and 2016–17, checking each against `c_agrgtd_tnovr_amt`.

A.7 Derived variable: the potential carryback

The carryback items on the tax return are populated only for firms that lodged a claim. To measure the refund available to a firm that did not claim -- which is the object of interest throughout the paper -- we reconstruct the statutory entitlement from the firm's own reported history. This subsection sets out that reconstruction.

A.7.1 Notation and inputs

For firm i and income year t, define the following, all measured after aggregation to the firm level and before deflation.

- TI_{it}, taxable income or loss (`c_taxinc`), which is negative in loss years.
- $L_{it} = \max\{0, -TI_{it}\}$, the tax loss of the year.
- $E_{it} = \max\{0, X_{it} - Z_{it}\}$, net exempt income, where X_{it} is exempt income (`c_exemptax`) and Z_{it} is non-deductible exempt income expenditure (`c_nondedex`).
- T_{it}, the income tax liability of the year: tax on taxable income (`c_taxti`), label T1 of the calculation statement.
- τ_{it}, the statutory rate constructed in Appendix A.6.
- F_{it}, the closing franking account balance (`c_div_frnkg_acnt_closg_bal_amt`), with `c_franbalc` used where the closing balance is missing.
- Ω_{it}, an indicator for statutory eligibility, equal to one if aggregated turnover (`c_agrgtd_tnovr_amt`) is below \$5 billion, the entity lodged a company return (`bit_comp_yyyy` $= 1$), and $F_{it} > 0$.

Three of these deserve comment. First, we use net exempt income rather than gross exempt income because the statute requires the loss carried back to a year to be reduced by that year's net exempt income to the extent it has not already been absorbed; the return reports exempt income and the associated non-deductible expenditure separately, and their difference is the relevant quantity. Second, we use tax on taxable income rather than tax payable as the income tax liability of the gain year. Tax payable is net of offsets already applied, and would understate the pool of prior-year tax available for refund. Third, the franking cap binds at the end of the claim year, not of the loss year, so a firm that paid large franked dividends between the profitable year and the claim will have less carryback capacity than its profit history alone would suggest. This is the source of variation exploited in Section 6.

A.7.2 The calculation

Table A.6 lists the admissible pairings of a loss year with an earlier profitable year, together with the order in which we apply the loss and the rate at which it is valued. Consider a firm claiming on its 2020-21 return in respect of a 2020-21 loss, the leading case in our analysis. Write $\ell_0 = L_{i,2021}$ for the loss, and let the ordered list of gain years be (2020, 2019), denoting the 2019-20 and 2018-19 income years by their end years. The loss is applied sequentially:

$$a_{i,2020} = \min\Big\{\max\{\ell_0 - E_{i,2020},\, 0\},\ \max\{TI_{i,2020},\, 0\}\Big\}, \tag{14}$$

$$\ell_1 = \max\Big\{\ell_0 - E_{i,2020} - a_{i,2020},\ 0\Big\}, \tag{15}$$

$$a_{i,2019} = \min\Big\{\max\{\ell_1 - E_{i,2019},\, 0\},\ \max\{TI_{i,2019},\, 0\}\Big\}. \tag{16}$$

That is, the loss is first reduced by the net exempt income of the year it is being carried back to, then applied against that year's taxable income up to the amount of that taxable income. Whatever survives is carried to the next year and the same operation is repeated. The potential refund is the sum of the amounts so applied, valued at the loss-year rate and capped year by year at the income tax liability actually incurred in the gain year, with a final cap at the franking account balance:

$$P_{i,2021} \;=\; \Omega_{i,2021} \cdot \min\left\{ \sum_{s \in \{2020,\,2019\}} \min\big\{\tau_{i,2021}\, a_{i,s},\ T_{i,s}\big\},\ \ F_{i,2021} \right\}. \tag{17}$$

Three features of equation (17) follow the statute rather than our discretion. The rate is that of the loss year, not of the year the loss is carried back to, so a 2020-21 loss carried into a year taxed at 27.5 percent is nonetheless refunded at the 2020-21 rate. The liability cap applies separately to each gain year, so a large loss cannot recover more from a given year than the tax that year generated. And the franking cap applies to the total, after summing across gain years.

The same construction is applied to the other claim years in Table A.6, extending the ordered list of gain years accordingly. Three accounting details arise once more than one loss year can be carried into the same gain year. First, a firm with losses in both 2019-20 and 2020-21 could claim in respect of both on its 2020-21 return; we apply the earlier loss year first, then reduce the remaining taxable income and remaining income tax liability of each gain year by the amounts already absorbed, so that no dollar of prior-year tax is refunded twice. Second, once a firm has actually claimed in an earlier year, the taxable income and liability of the gain years are depleted for the purposes of any subsequent claim; we carry this depletion forward using observed claim amounts where a claim was made and our own imputed amounts otherwise. Third, a loss cannot be carried back if it has

already been deducted against income in an intervening year, so we net `c_ddctd_tl_amt` out of the loss pool before applying the calculation, and check the residual against tax losses carried forward, `c_tlcf_amt`.

A.7.3 Substituted accounting periods

A small number of companies do not balance on 30 June. For these early and late balancers the return does not carry the year-specific carryback labels but a generic set instead: the loss chosen to be carried back to the earliest year, the corresponding amount for any middle year, and the net exempt income and income tax liability of the earliest year. These are the four early balancer items listed in Panel G of Table A.3. We apply the same algorithm to these firms, using the balancing date reported on the return to map their accounting period to our financial-year convention, and we flag them so that the results can be re-estimated with them excluded.

A.7.4 Validation against observed claims

The reconstruction can be checked directly on the subsample of firms that did claim, for whom the true entitlement is reported. Regressing the constructed potential carryback $P_{i,2021}$ on the refund actually claimed, `c_lss_crry_bck_to_amt`, among claimants yields a coefficient of 0.988, with a standard error of 0.164 and an R^2 of 0.977. The near-unit slope and very high fit reflect two distinct facts. The behavioral fact is that take-up operates almost entirely on the extensive margin: a firm that claims at all claims essentially the whole of what it is entitled to, so that the observed refund and the statutory entitlement coincide for this group. The measurement fact, which is what justifies using the constructed variable for the much larger population of non-claimants, is that our implementation of the statutory formula reproduces the entitlement itself. The residual variation is concentrated among firms with substituted accounting periods and among consolidated groups whose head company reports on behalf of subsidiaries.

Table A.6: Admissible loss-year to gain-year pairings and the order of application

Claim year (return lodged)	Loss year Y_L	Gain years Y_G, in order of application	Rate applied
2020–21	2019–20	2018–19	$\tau_{i,2019–20}$
	2020–21	2019–20, then 2018–19	$\tau_{i,2020–21}$
2021–22	2021–22	2020–21, then 2019–20, then 2018–19	$\tau_{i,2021–22}$
2022–23	2022–23	2021–22, then 2020–21, then 2019–20, then 2018–19	$\tau_{i,2022–23}$

Notes: A firm lodging its 2020–21 return could claim in respect of a 2019–20 loss, a 2020–21 loss, or both; losses incurred in 2019–20 could not be claimed until the 2020–21 return was lodged. The statute imposes no ordering requirement across gain years, but applying the loss to the most recent profitable year first weakly maximizes the refund because the income tax liability of nearer years is more likely to remain unabsorbed, and it matches the order in which claims are in fact reported. When more than one loss year is carried back to the same gain year we apply the earlier loss year first and reduce the remaining taxable income and income tax liability of that gain year accordingly, so that no dollar of prior-year tax is refunded twice. Each pairing corresponds to a distinct BLADE variable of the form `c_fy_`Y_L`_to_fy_`Y_G`_tax_lss_cba`, where years are labeled by their end year.

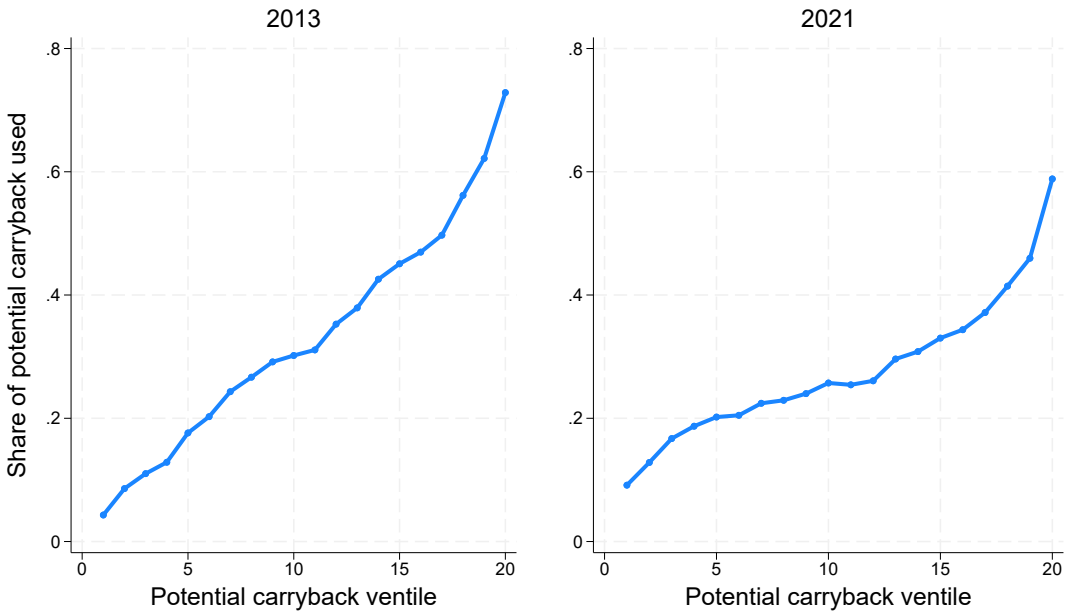
A.8 Sample selection

The employment and revenue screens of Section 4.2 are applied at the level of the firm rather than the firm-year: a firm that satisfies the condition in any pair of consecutive years is retained for its entire spell in the panel. This matters because the alternative -- screening firm-years -- would drop

exactly the observations in which a firm's employment or revenue collapsed, which is the outcome we are trying to measure. Employment is headcount from the withholding module, supplemented from 2019-20 by the count of employees paid in the monthly payroll module; where the two disagree we use the withholding measure for the annual screens and the payroll measure for the monthly analysis. Revenue is total sales from the activity statement module, deflated to 2023 dollars.

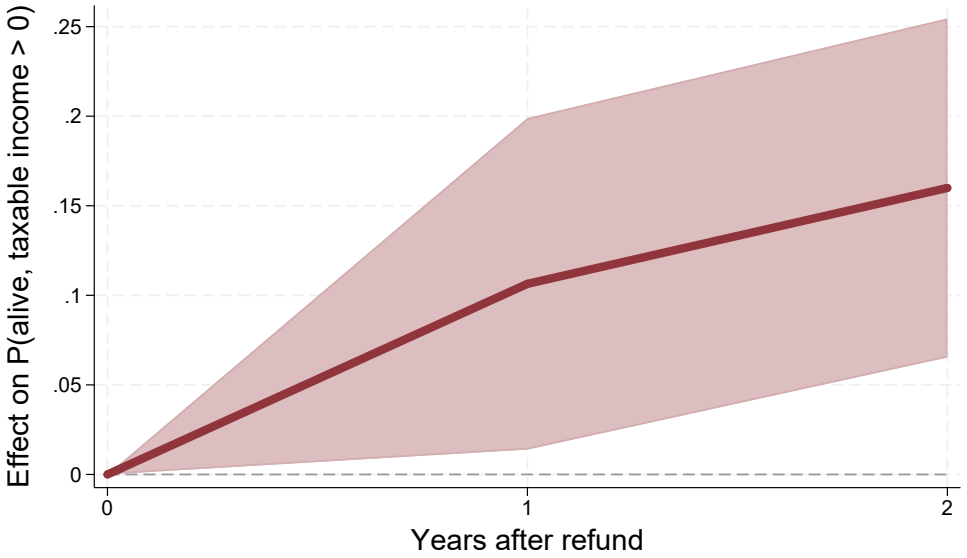
B Additional Figures and Tables

Figure B.1: Share of potential refund dollars used, 2012–13 and 2020–21 episodes



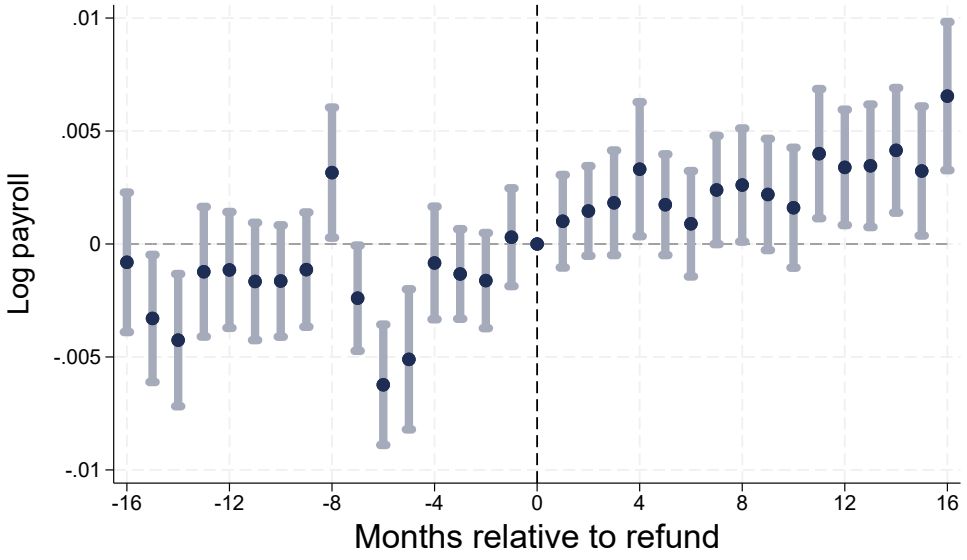
Notes: Firms eligible for a carryback are ranked into ventiles of the potential carryback, separately for each episode. The figure plots the mean share of the potential carryback actually claimed within each ventile, for the 2012–13 episode (left panel) and the 2020–21 episode (right panel). Statistics are rounded in accordance with ABS disclosure rules.

Figure B.2: The effect of carrybacks on survival with positive taxable income



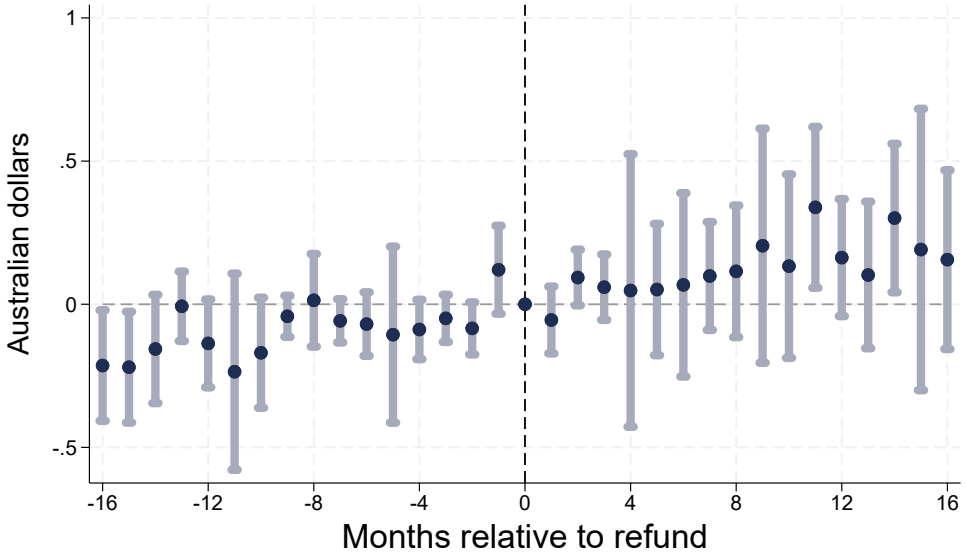
Notes: The figure plots the coefficients β_h from equation (12) with the outcome equal to one only if the firm is alive and reports positive taxable income at horizon h, with 95 percent confidence intervals. Statistics are rounded in accordance with ABS disclosure rules.

Figure B.3: The effect of carrybacks on monthly payroll: log specification



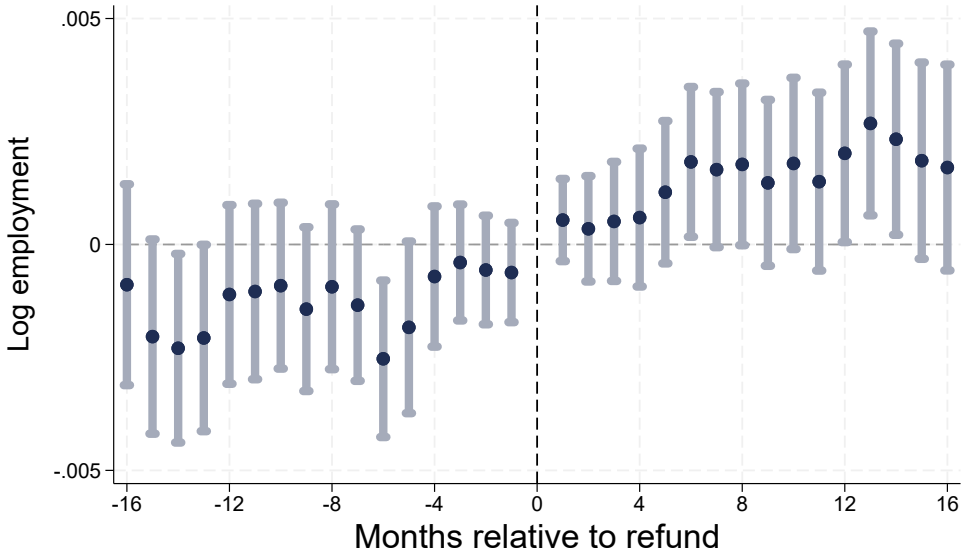
Notes: The figure plots the differential path β_k from equation (13) with the logarithm of total monthly payroll as the outcome, for event months $k \in \{-16, \ldots, 16\}$ relative to the expected refund, with 95 percent confidence intervals. Statistics are rounded in accordance with ABS disclosure rules.

Figure B.4: The effect of carrybacks on monthly payroll: entropy reweighting



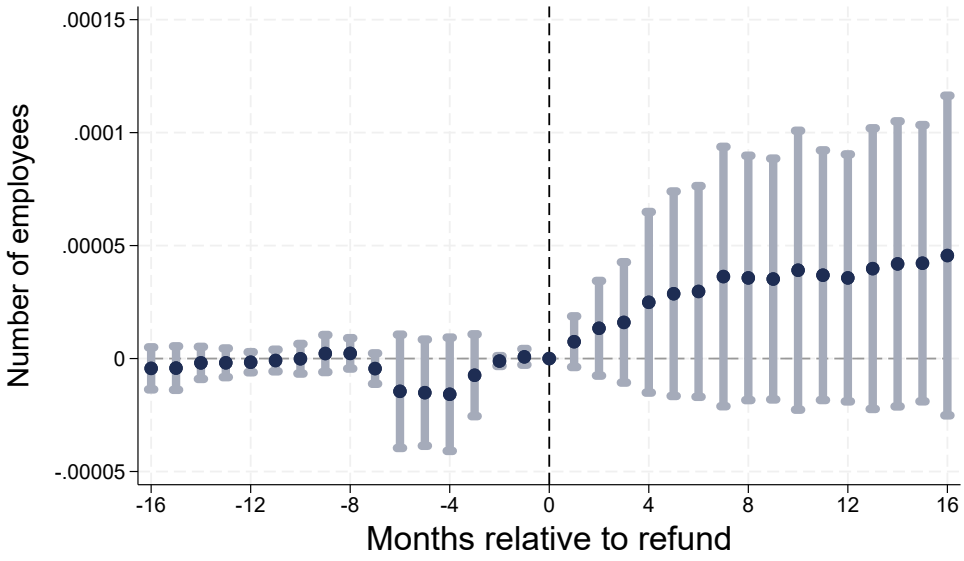
Notes: The figure plots the differential path β_k from equation (13) with total monthly payroll in dollars as the outcome, for event months $k \in \{-16, \ldots, 16\}$ relative to the expected refund, with 95 percent confidence intervals. The pass-through population is reweighted toward the corporate covariate means by entropy balancing in place of the matched pairing. Statistics are rounded in accordance with ABS disclosure rules.

Figure B.5: The effect of carrybacks on employment: log specification



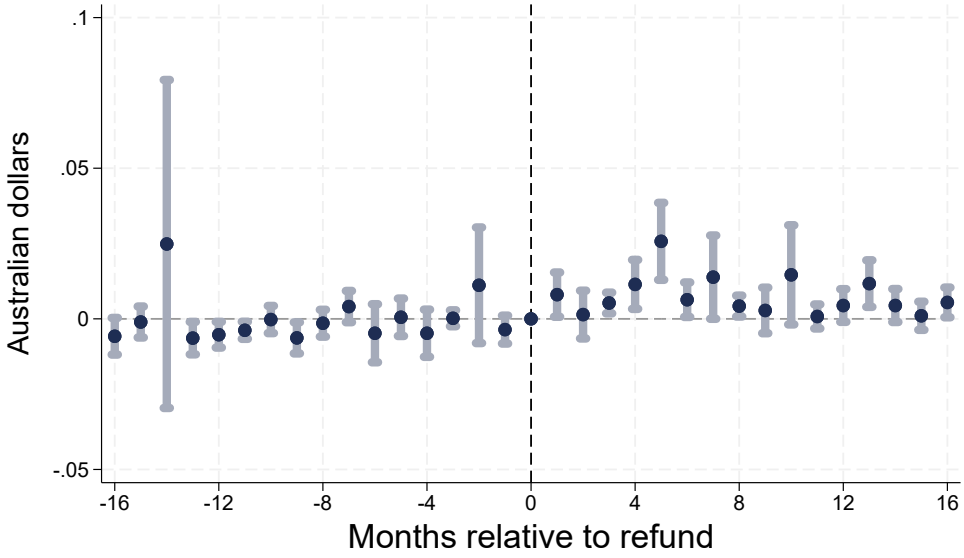
Notes: The figure plots the differential path β_k from equation (13) with the logarithm of employment as the outcome, for event months $k \in \{-16, \ldots, 16\}$ relative to the expected refund, with 95 percent confidence intervals. Statistics are rounded in accordance with ABS disclosure rules.

Figure B.6: The effect of carrybacks on employment: entropy reweighting



Notes: The figure plots the differential path β_k from equation (13) with the number of employees paid as the outcome, for event months $k \in \{-16, \ldots, 16\}$ relative to the expected refund, with 95 percent confidence intervals. The pass-through population is reweighted toward the corporate covariate means by entropy balancing in place of the matched pairing. Statistics are rounded in accordance with ABS disclosure rules.

Figure B.7: Payroll paid to new hires, on its own vertical scale



Notes: The figure plots the differential path β_k from equation (13) with payroll paid to new hires as the outcome, on its own vertical scale rather than the common scale of Figure 10, for event months $k \in \{-16, \ldots, 16\}$ relative to the expected refund, with 95 percent confidence intervals. Statistics are rounded in accordance with ABS disclosure rules.