Submission: Reducing barriers to business dynamism

Prosperity Institute, Mark Humphery-Jenner

1 About the author

The prosperity institute is dedicated to improving Australian prosperity and increasing individual freedom and liberty. It is not associated with any political party or advocacy organization

Mark Humphery-Jenner operates Prosperity Institute and is the author of this submission. He founded Otso Capital, is an Associate Professor of Finance at UNSW, and sits on the Management Committee of Sydney Angels and Investment Advisory Committee of Australian Vision Research. He has written more than 40 peer reviewed articles in finance and law journals.

2 High CGT rates undermine business dynamism and investment

This submission is primarily concerned with recent changes to tax rules. We highlight that rigorous analysis shows that these tax rules will both reduce total government revenue and will reduce overall welfare. Our firm recommendation is that capital gains tax be significantly reduced, tending towards zero. Our modeling shows that such a reduction would help to maximize government revenue and maximize welfare. The full modeling is in the attached appendix.

Revenue-maximizing rates are high only under restrictive assumptions. Allowing for elastic investment, fiscal spillovers, growth effects, substitution toward lower-taxed assets or jurisdictions, and realization-based lock-in materially lowers optimal rates. Empirically grounded calibrations imply revenue-maximizing rates are low, often near zero, and below common statutory rates; welfare-maximizing rates are lower still.

High CGT rates undermine business dynamism in several ways, as we prove in our extended model:

1. The principal place of residence (PPR) is now Australia's best investment. However, investors also see a clear low taxed domestic alternative: owner occupied principal place of residence (PPR). The PPR has no tax. A mortgage interest rate is currently around 6%. This is now equivalent to earning 12% before tax risk free on another investment. But, we must then add liquidity and risk premiums for stock, and especially startup, investments. This means that the high CGT rates make investing in dynamic businesses less attractive.
2. High CGT rates deter investment in primary shares, such as in startups. We are aware that there are moves to provide some carve outs to recent CGT increases. However, they are complex, discretionary, and uncertain. They are insufficient to offset the recent CGT increases. Furthermore, they cap-out for larger companies, making them irrelevant for investment sin late stage companies, which can also generate growth.
3. High CGT rates undermine the secondary market. The preponderance of evidence shows that a liquid and functional secondary market is essential for a functioning primary market. For example, a secondary market helps firms to properly price, and sell, shares to investors. Thus, where CGT rates undermine the secondary market, they also interfere with the primary market.
4. High CGT rates create a lock-in effect, which prevents capital going to its best use: Specifically, a high CGT rate deters investors from selling shares in a low growth asset to reallocate to a high growth asset as the new high growth asset may take years to recapture even the initial tax outlay.
5. High CGT rates deter founders from moving to Australia and can encourage founders to leave. We are aware of dismissive talk that founders might not leave. They will simply accept high CGT because the weather is nice in Australia or because of inertia. But, such dismissiveness ignores the flip-side: fewer founders will come to Australia. For example, whereas Australia could attract a NZ founder with

sufficiently low CGT, that would now not be the case. The high CGT rates reduce the pool of founders relative to what would otherwise be the case.

6. Capital investment and gains are not income. The analogy between the two is as fallacious as saying that because apples grow on trees, so too must carrots. The two are separate activities. A simple illustration proves the point: if a company fails, the government will not refund the capital deployed to establish it, much of which will be from after-tax sources. Capital related gains have a significantly different risk profile from wages. And, as we illustrate below, a revenue-maximizing capital gains rate is low, meaning that an income tax cut would be more available were CGT rates to be lower.
7. High CGT rates and franked dividends: Fully franked dividends are now relatively more attractive, especially if invests have stable liquidity needs. This is because the tax rate on sales is 47% of the real gains under the new rules, for a person on the top tax rate. By contrast, the tax rate on fully franked dividends is near 17%. Thus, unless an investor can for sure compound gains, which are taxed annually at the corporate rate of 30%, the investor would prefer fully franked dividends.

For these reasons, there is significant academic evidence that CGT rates have a material impact on the venture capital ecosystem, as we detail in the appendix. We therefore recommend that a government that is serious about business dynamism significantly reduce CGT, with optimal rates being near zero.

3 Regulatory impost and regulatory litigiousness

Regulatory factors significantly reduce business dynamism. A recent case-in-point is the incoming AML regulations which impose additional obligations on myriad professional services firms. This is not an isolated incident. Policy makers are too quick to impose regulations that other people must pay for. In turn, these regulations increase startup costs and entrench incumbents.

Overarching recommendations are:

The productivity commission should have regard to the following:

1. **Regulators do not bear the cost of failed enforcement actions**. For example, should a regulator, such as ASIC or the ATO, pursue a failed case, the people who decided on the case suffer no penalty, no matter how doomed the case was. This is a problem because such regulatory activities cost the business money and force the business to devote more resources to professional services. We recommend a penalty system whereby cases of sufficiently bad quality result in penalties for the regulator.
2. **Burden of proof**: in all cases, the burden of proof should fall on the regulator to prove allegations. This is the case whether it be ASIC, APR, FWO, or ATO. The fact that the burden falls on businesses in some cases means that the business must spend additional time and resources just to deal with regulatory matters. All cases brought must cross a sufficient threshold of evidence so as to avoid businesses expending money on regulatory compliance.
3. **Audit of regulations**: Australia must resist the urge to increase regulation. This applies across the field, including in the area of environmental regulation. For example, requiring Scope 3 emissions data imposes a cost on business. Recent AML changes are another example: they impose additional costs, and slow business activity, but provide no benefit to the firms or to the clients. All such additional costs should be shelved. Australia must also consider a system whereby businesses that are impacted by regulatory costs can effectively have a say on such costs. We must also audit regulations systematically to reduce and eliminate regulations that are not fit for purpose.

Appendix: Optimal Capital Gains Tax as an Incentive Contract[*]

Mark Humphery-Jenner[†]

Abstract

We recast optimal capital gains taxation as a moral-hazard problem: the government is a principal, the investor an agent whose unobservable investment-effort produces a noisy gain. The tax rate is the slope of a linear sharing rule. We distinguish revenue-maximizing from welfare-maximizing rates, reflecting the relative value of public-versus-private funds. Revenue-maximizing rates are high only under restrictive assumptions. Allowing for elastic investment, fiscal spillovers, growth effects, substitution toward lower-taxed assets or jurisdictions, and realization-based lock-in materially lowers optimal rates. Empirically grounded calibrations imply revenue-maximizing rates are low, often near zero, and below common statutory rates; welfare-maximizing rates are lower still.

[*] This paper benefited from feedback received, including on earlier versions and policy discussions, at the Senate Inquiry into the Capital Gains Tax Discount (February 2026), Wilson Asset Management Investor Discussion Group (June 2026). We are thankful for feedback on this, and earlier drafts and related analysis from Geoff Wilson and Peter Swan.
[†] UNSW Business School and Otso Capital. Corresponding author: MLHJ@UNSW.EDU.AU

1 Introduction

How heavily should a government tax capital gains? The question is old, and the purported answers in practice are startlingly dispersed. At least amongst policy makers. Across advanced economies the headline rate on realized gains ranges from zero to roughly half, and *within* a single jurisdiction the same dollar of gain can face wildly different treatment depending on the asset that produced it, the vehicle that held it, and where the investor happens to reside. This dispersion is hard to reconcile with a settled theory of what the rate *should* be.

The empirical evidence shows that capital gains tax (CGT) rates can have a significant impact on economic activity, including on innovation. For example, Dimitrova and Eswar (2023) show that higher CGT rates significantly reduce venture capital (VC) investment. Edwards and Todtenhaupt (2020) find that reducing CGT can spur significant investment in qualifying investments; in their case, in startups. Given the well-documented link between VC and innovation, higher CGT would thus logically result in less innovation and growth through the venture and startup channel. Similarly, Akcigit et al (2016) show that inventors, especially superstar inventors, will shift locations in response to tax rates. The real effects of CGT rates extend further: for example, Ayers et al (2003) document that high CGT rates influence acquisition premiums, implicitly interfering with the market for corporate control and its attendant growth and productivity effects. In turn, the empirical literature raises the question of what precisely that optimal capital gains tax might be.

The dominant theoretical frameworks approach the problem through the lens of private information about *types* rather than about encouraging people to exert effort through human or financial capital. The Ramsey tradition studies a planner choosing linear taxes to finance spending in a representative-agent economy. Seminal papers in this paradigm famously conclude that the optimal long-run capital tax is zero (Chamley, 1986; Judd, 1985). However, subsequent critiques have suggested this result is sensitive to the model's assumptions (Straub and Werning, 2020). The Mirrleesian tradition treats the government as a principal screening agents who differ in *unobservable productivity*, and derives capital taxation as an intertemporal wedge whose role is to relax the incentive constraints of high types (Golosov et al., 2003; Kocherlakota, 2005; Mirrlees, 1971). The two traditions disagree about the answer, but they share a premise: the binding friction is hidden *type*, and the capital tax is an instrument in a redistributive screening problem.

This paper analyzes a different friction: how to incentivize people to deploy capital. We contend that the challenge in capital markets is arguably not that the government cannot observe who is talented, but that it cannot observe what investors *do*. An investor chooses how much "effort", broadly defined, to exert. Such effort includes capital deployment, search, due diligence, monitoring, and continued analysis. This effort, which includes mere capital deployment, produces a return. The investor's effort is both unobservable and produces a noisy payoff. This is akin to the canonical moral hazard problem documented in Holmstrom (1979). Here, the government is the principal; the investor is the agent; the realized capital gain is a noisy signal of unobservable effort; and the capital gains tax is the *sharing rule* through which the principal splits output with the agent. Optimal capital taxation becomes the design of an incentive contract.

The reframing earns its keep in two ways. First, it speaks to tax *structure*, not just the rate. Holmstrom and Milgrom (1987) show that the optimal contract in this environment is linear in realized output. This reflects the practical operation of capital gains taxes context: In most countries, CGT rates are flat and linear. Even in countries where rates scale with the individual's income, the bulk of revenue will derive from high-income earners on the top marginal rate. Second, the central tradeoff is transparent and familiar. A higher tax is a lower-powered contract; it dampens unobservable effort exactly as weak incentives dampen an employee's. Whereas the screening tradition finds capital taxation *desirable* because it slackens the incentive constraints of the rich, the moral-hazard framing recovers the standard agency distortion. That is, capital taxation dampens effort, which influences total welfare and revenue collected.

We build the argument cumulatively. We begin with the baseline Holmstrom and Milgrom (1987) model applied to risky investing. Effort is linear in the retained share, so the tax base is unit-elastic in the net-of-tax rate, and the static revenue-maximizing rate is one half. However, we then show that one half is myopic. Investment is not a one-time lottery: it throws off fiscal spillovers. These include corporate tax, payroll tax, employment (and associated income tax). Once the government credits this ancillary revenue, the revenue-maximizing rate falls below one half and the welfare-maximizing rate falls further. Beyond a threshold spillover the optimal policy is not a tax but a *subsidy*, which rationalizes targeted credits such as Australia's ESVCLP and ESIC regimes as the Pigouvian correction implied by the contract. This effect is even greater if we also consider the impact of investment on long term economic growth, which the empirical literature shows to be the case (Benedetti-Fasil et al., 2022; Lindsey and Stein, 2026). In this case, the optimal capital tax falls further.

The situation becomes more stark when investors have multiple investment alternatives with different tax rates. If the agent were to divert effort to a lower-taxed or untaxed alternative, we face a multitask setting, analogous to that in Holmstrom and Milgrom (1991). This could arise if different asset classes are taxed differently or capital is mobile. In this case, the optimal rate on the taxed channel falls in the degree of substitutability and approaches zero as the alternative becomes a perfect substitute. This delivers a rationale for taxing like assets uniformly. Finally, crediting the present value of the future tax stream that investment sets in motion pushes the optimal rate down once more.

Across every extension the conclusion is the same: the optimal capital gains rate lies below the static revenue-maximizing rate that ignores ancillary tax-revenue that derives from investment. The welfare-maximizing tax rate is lower still. This is notable. It supports the idea that capital should not simply be taxed like labor, as some have contended (see e.g., Stewart, 2026). Rather, capital ought be taxed optimally in order to encourage its deployment. Furthermore, our results emphasize that increasing capital gains tax rates does not necessarily increase capital gains tax revenue, and especially can reduce total welfare.

Our reframing is closest in spirit to, but functionally distinct from, literature that places risk and effort inside the tax problem. The taxation-of-risk tradition asks how proportional taxes with loss offsets shape portfolio risk-taking (Domar and Musgrave, 1944), and a related strand studies the taxation of entrepreneurial risk-taking and entry (Cullen and Gordon, 2007; Gentry and Hubbard, 2000; Gordon and Sarada, 2018; Keuschnigg, 2004; Keuschnigg and Nielsen, 2004, 2003; Scheuer, 2014). These literatures share our premise that tax shapes costly, risky private decisions. They differ in what the rate *is*. We interpret the capital gains rate as the incentive intensity of a linear sharing rule between a government-principal and an investor-agent whose effort is unobservable. The tax rate is then the slope of the contract, rather than an intertemporal wedge or a portfolio-risk parameter. To our knowledge we are the first to read the rate this way, and the reading is not merely terminological: it is what generates our distinctive results. The sharing-rule lens delivers (i) a fiscal-spillover threshold beyond which the optimal contract credits, rather than taxes, investment, rationalizing targeted regimes such as Australia's ESVCLP and ESIC as the Pigouvian correction the contract implies; (ii) a multitask neutrality result under which taxing like assets uniformly is precisely what forecloses base-shifting; and (iii) an interaction between realization-based timing and the growth dividend, whereby a higher rate raises the reallocation hurdle and suppresses the very activity that produces downstream revenue.

None of these follows from a hidden-type screening problem, and none is visible when the rate is treated as a portfolio parameter rather than a contract slope.

Our paper makes a significant contribution to the optimal taxation literature. While our results do relate to existing capital tax frameworks, we propose a different approach to considering optimal capital taxation. Relative to Chamley-Judd style results, we obtain a low (and sometimes negative) optimal capital tax. However, we obtain such a result via a different mechanism. Having reached a similar end by discrepant means, to paraphrase Montaigne (1580), we help to cross-validate the logic in Chamley-Judd. Our paper contrasts to the Mirrleesian literature. That literature asks how capital taxes should redistribute across heterogeneous types, whereas we ask how they should be set when investment effort is hidden and elastic.

Our paper also speaks to the literature on behavioral responses to tax rates. An extant stream of literature analyzes how high tax rates might go before the government loses revenue, or economic activity declines. This manifests in the literature on the elasticity of taxable income (Feldstein, 1999, 1995), and the approach to deriving tax rates from estimated behavioral responses (Chetty, 2009; Saez, 2001). That literature focuses on how economic activity alters when tax rates change. We capture such an elasticity from a moral hazard perspective. Whereas the extant approach treats elasticity as a parameter to be measured, we also incorporate such elasticity effects explicitly into our model. For example, we explicitly incorporate growth effects, or the decision to shift activity into lower taxed activities or jurisdictions. We also explicitly consider the asset lock-in effect. In so doing, we help to unpack direct channels through which the elasticity of behavior might manifest.

Finally, by modeling the investor as an agent who chooses effort and can found a firm, the paper speaks to the taxation of entrepreneurship and risk-taking (Cullen and Gordon, 2007;

Scheuer, 2014), which documents that tax reductions can stimulate entrepreneurship and economic growth (Sedlacek and Sterk, 2019). Our paper's theoretical results also help to cross-validate empirical findings that investors, especially superstar inventors, can be incentivized to move to lower tax jurisdictions (Akcigit et al., 2016).

The paper has significant policy implications. Some countries have looked to increase capital gains tax rates, believing that it will raise additional revenue. An example is Australia, which has moved move toward taxing capital at the top marginal tax rate of 47%.[1] Similarly, the United Kingdom increased its capital gains tax rate to 24% from 20%[2] and curtailed territorial taxation, and reduced tax free allowances. However, capital gains tax receipts decreased contemporaneously, falling between 10% (Barnett, 2025) and 18% (Agyemang, 2025), depending on the time period analyzed. This has anecdotally been attributed to factors such investors holding assets longer due to the reduction in tax free allowances (i.e., a lock-in effect), and high net worth individuals leaving the UK. However, investors also anecdotally moved to liquidate assets before tax hikes would take effect. The net result is that increasing tax rates can have behavioral effects that must be considered. Our results thus highlight that many advanced economies could be moving in the wrong direction if the goal is to maximize either revenue or welfare.

The remainder of the paper proceeds as follows. Section 2 develops the baseline model and the naïve revenue- and welfare-maximizing rates. Section 3 introduces fiscal spillovers and the subsidy threshold. Section 4 adds differently taxed and untaxed alternatives, and the now-versus-later consumption margin. Section 5 incorporates the growth dividend from investment. Section 6 explores the impact of lock-in effects owing to capital taxation. Section 7 applies our model to explore how our model functions in a real-world environment

[1] This is a function of the Treasury Laws Amendment (Tax Reform No. 1) Bill 2026, which would tax capital gains at the individual's marginal tax rate. The top marginal tax rate is 47%. The tax would be imposed on real gains (which, we will note, is what the present paper's tax rates also apply to).
[2] See Finance Act 2025 (UK), Section 7.

and what it means for real world revenue-maximizing and welfare-maximizing rates. Section 8 concludes.

2 The naïve first-pass approach to CGT

We start with the baseline Holmstrom and Milgrom (1987) model as applied to risky investing. We use this as a baseline from which to analyze the relationship between taxation and investment effort. Throughout this paper, returns are in *real* (not *nominal*) terms. We build upon this model in subsequent sections in order to make the model more realistic and incorporate the impact of investing on economic growth and taxation. The initial set up is as in Holmstrom and Milgrom (1987) applied mutatis mutandis to investing.

We note that the baseline model has several limits, which we relax in subsequent sections: (i) to begin with, the assumption of a quadratic cost function assumes a unit elasticity of investment-effort to tax rates (and to the keep rate, being $\beta = 1 - \tau$), (ii) the baseline model assumes symmetrical treatment of losses and gains, which may only be realistic for institutional investors, especially those who are profitable on net and can offset losses against gains, (iii) the baseline model ignores externalities from investing, such as corporate value creation and employment and the ancillary revenue from such activity, (iv) the baseline model ignores the possibility that an investor might shift capital and effort into consumption or a lower taxed alternative, (v) the baseline model ignores the economic-growth impact of investing, including the tax revenue that such growth adds, and (vi) the baseline model does not reflect investment lock-in and the impact this might have on returns, fiscal revenue, and growth.

2.1 Baseline Setup

The initial setup is straightforward. We have an initial investor. The investor chooses an unobservable amount of effort $e \geq 0$ to apply to capital investing. The model applies mutatis mutandis to the decision to found a company as the payoff to founding a company comes from capital gains. The effort encompasses capital deployment, analysis, and monitoring. The effort produces a risky gain. That is, the gain scales with the amount of capital and skill deployed.

$$Y = Ae + \epsilon, \qquad \epsilon \sim N(0, \sigma^2) \tag{1}$$

Where, $A > 0$ is the productivity of investment-effort (i.e., the percentage returns the investment produces) and ϵ represents the randomness of returns. Thus, Y can be thought of as the dollar return to that investment effort.

The government levies a proportional capital gains tax τ on the realized capital gains. The investor keeps $\beta Y = (1 - \tau)Y$. This fits the CGT context directly, where capital gains tax rates are often a constant, linear, percentage rate. And, even where capital gains are taxed at the investor's marginal tax rate, we note that the majority of revenue raised will be from high net worth individuals, making the tax rate effectively linear. We also note that the tax-structure both matches the empirically observed tax structure (except for taxes often being on nominal gains), and reflects the arguments in Aguiar et al (2024), who argue that optimal taxes are based on real gains, are only on realized gains, not unrealized changes in asset prices, are on gains (rather than gross revenue), and that investors are "compensated" for realized losses.

Exerting effort is costly. We initially assume a quadratic cost function with $C(e) = \frac{k}{2}e^2$ with $k > 0$. The assumption of quadratic cost significantly influences the optimal tax rate. We address this assumption below. In our context, this costly effort represents the fact that the investment itself is costly and the investor must exert effort to select assets. We assume that the investor has constant absolute risk aversion (CARA), as is standard in the literature. We impose a risk aversion parameter, $a > 0$, over the risky payoff $x = \beta Y - C(e)$. Thus, the investor's utility is:

$$U(x) = -\exp(-ax)$$

The investor will invest only if the investment produces a return at least as high as their outside option. That is, if $U \geq \overline{U}$. The government's problem is to choose τ based on the investor's anticipated response. By contrast, if the resultant payoff is below the reservation rate, the investor might simply consume (more on that below) or leave for a more tax-efficient environment. Such capital flight is well-documented, and is especially poignant for high net worth individuals whose income derives from investing or who are mobile. For example, Kleven et al (2013) show that sporting superstars are attracted to low tax environments. Similarly, investors, especially foreign or superstar inventors, are attracted to low tax environments (Akcigit et al., 2016). While we do not explicitly model the differential impact of taxes on differently talented individuals, the participation constraint appears to especially bite for the types of individuals a country might most want to retain. That is, superstar investors, which presumably drive supernormal growth relative to standard inventors, are the most likely to move. This model focuses on homogeneous agents, rather than exploring agents of different types. However, the intuition from the empirical literature on mobility grounds the role of the participation constraint: the participation constraint binds when the outside option is attractive relative to the after tax return.

2.1.1 The investor's choice

The investor chooses the optimal amount of effort with respect to their anticipated payoff. As is well established, this is equivalent to maximizing the certainty equivalent of the payoff. In order to make this paper self-contained, we establish this as Lemma 1; however, this is a well-established result.

Lemma 1 (Certainty equivalent of a CARA-normal payoff). If $x \sim N(\mu, v)$ and $U(x) = -\exp(-ax)$, then $E[U(x)] = -\exp\left(-a\left[\mu - \frac{a}{2}v\right]\right)$. Maximizing expected utility is therefore equivalent to maximizing the certainty equivalent $\mu - \frac{a}{2}v$.

We therefore apply Lemma 1 to $x = \beta Y - C(e) = \beta(Ae + \epsilon) - \frac{k}{2}e^2$. The mean and variance are $\mu = \beta Ae - \frac{k}{2}e^2$ and $v = \beta^2\sigma^2$. Therefore, the certainty equivalent is:

$$CE(e) = \underbrace{\beta Ae}_{\text{Exp. retained gain}} - \overbrace{\frac{k}{2}e^2}^{\text{Effort cost}} - \underbrace{\frac{a}{2}\beta^2\sigma^2}_{\text{Risk premium}} \tag{2}$$

The investor chooses the optimal effort e so as to maximize the certainty equivalent. The first order condition $\frac{d(CE)}{de} = \beta A - ke = 0$ gives:

$$e^* = \frac{\beta A}{k} = \frac{(1-\tau)A}{k} \tag{3}$$

Therefore, effort is linear in the retained share. If the government increases capital tax, then investors deploy proportionately less. The intuition is straightforward: if there is less payoff to investing, then investors will devote their human and financial capital elsewhere. We also note that the certainty equivalent is then as follows:

$$CE^* = \frac{\beta^2 A^2}{2k} - \frac{a}{2}\beta^2\sigma^2 = \frac{\beta^2}{2}\left(\frac{A^2}{k} - a\sigma^2\right)$$

Recall that the investor participates only if the payoff is at least as good as the outside option. That is, only if $CE^* \geq \bar{u}$, where $\bar{u}$ denotes the value of that outside option. Therefore, the investor invests only if:

$$\frac{(1-\tau)^2}{2}\left(\frac{A^2}{k} - a\sigma^2\right) \geq \bar{u}$$

Therefore, there is a tax threshold τ at which the investor exits entirely. For example, they might prefer to deploy their scarce resources on leisure or simply move overseas. We denote this threshold tax as $\bar{\tau}_{\text{participation}}$. Thus, the investor participates only if:

$$\tau \leq \bar{\tau}_{\text{participation}} = 1 - \sqrt{\frac{2\bar{u}}{\frac{A^2}{k} - a\sigma^2}}$$

We capture this in the following Lemma:

Lemma 2 (Investor effort and participation, second order conditions in Appendix A.1). Facing a proportional tax τ with retention $\beta = 1 - \tau$, the investor's unique optimal effort is $e^*(\tau) = \frac{\beta A}{k}$ with certainty equivalent $CE^*(\tau) = \frac{\beta^2}{2}\left(\frac{A^2}{k} - a\sigma^2\right)$. Effort is linear in the retained share, and the expected base $E[Y] = Ae^*$ has constant elasticity of one with respect to the net-of-tax rate (i.e., keep rate) β. The investor participates if and only if $\tau \leq \bar{\tau}_{\text{participation}} = 1 - \sqrt{\frac{2\bar{u}}{\frac{A^2}{k} - a\sigma^2}}$.

2.1.2 The government's decision

The government decides on the tax rate. For present purposes, we assume that the government aims to either maximize revenue or welfare. It is a value-judgment as to whether the government *should* attempt to maximize revenue or whether the government should attempt to minimize the amount of tax obtained. The direct capital gains revenue is simply the tax rate multiplied by the expected gain:

$$R(\tau) = \tau E[Y] = \tau(Ae^*) = \frac{\tau(1-\tau)A^2}{k}$$

Differentiating with respect to τ gives $R'(\tau) = \frac{A^2}{k}(1-2\tau)$. So, the first order condition gives:

$$\tau_{Naive}^{R} = \frac{1}{2} \tag{4}$$

However, as will become apparent, this is myopic. It ignores the impact of investing on economic growth, which in turn generates future tax revenue. Phrased differently, this is a static point-in-time revenue maximum. It is not the revenue maximum after incorporating growth effects. We address this in following sections.

We next distinguish between the revenue maximizing rate and the welfare maximizing rate. This is important because maximizing government revenue can destroy private surplus. The total welfare is a weighted sum of the private surplus and the government's revenue. To capture this, we use a parameter λ that represents the relative social value of a public dollar. Notably, this public revenue pertains to the public use of money and does not reflect any rebate to the investor (i.e., were the tax rate to be negative). Thus, intuitively, the tax revenue R funds a public good valued at λ per dollar. This is akin to the logic presented in Hendren and Sprung-Keyser (2020) in relation to the relative social value of a public dollar. Naturally, the value of that parameter will depend on government efficiency and corruption.

$$W(\tau) = \underbrace{CE^*(\tau)}_{\text{Private surplus}} + \lambda \underbrace{R(\tau)}_{\text{Government revenue}} \tag{5}$$

We can now explore the tradeoff between private surplus and government revenue. Let us define $m = \frac{A^2}{k}$ and $s = \frac{a\sigma^2}{m}$. Then, $CE^* = \frac{\beta^2}{2}m(1-s)$. This is a rearrangement of the foregoing value of CE^*. Therefore,

$$W(\tau) = \frac{(1-\tau)^2}{2}m(1-s) + \lambda\tau(1-\tau)m$$

We can now differentiate and divide by m to obtain:

$$\frac{W'(\tau)}{m} = -(1-\tau)(1-s) + \lambda(1-2\tau) = 0 \ \rightarrow \tau[(1-s) - 2\lambda] = (1-s) - \lambda$$

Solving this, and noting that $s = \frac{a\sigma^2}{m}$ and $m = \frac{A^2}{k}$ gives the following where $s = 0$:

$$\tau^*_{\text{Naive, Welfare}} = \frac{\lambda - 1}{2\lambda - 1} \tag{6}$$

This has a very clear interpretation. We plot the welfare maximizing rate in Figure 1. If the government values private and government capital equally, then $\lambda = 1$ and the optimal tax rate τ^* tends towards zero. This is because the government would be content with the individual retaining all their assets. By contrast, as $\lambda \rightarrow \infty$. Then $\tau^* \rightarrow \frac{1}{2}$. That is, we only get a 50% tax rate if the government values government revenue infinitely more than private capital. It is a normative question about whether such a value-judgment is appropriate. We capture this in the following proposition, below.

[Insert Figure 1 about here]

Proposition 1 (Baseline optimal rates). In the single-asset economy with $W(\tau) = CE^*(\tau) + \lambda R(\tau)$: (i) the revenue-maximizing rate is $\tau^R = \frac{1}{2}$; (ii) the welfare-maximizing rate is $\tau^* = \frac{\lambda+s-1}{2\lambda+s-1}$, where $s = \frac{a\sigma^2 k}{A^2}$. In cases where $s = 0$, this is $\tau^* = \frac{\lambda-1}{2\lambda-1}$; (iii) τ^* is increasing in λ; with $s = 0$, $\tau^* \to 0$ as $\lambda \to 1$, and $\tau^* \to \frac{1}{2}$ as $\lambda \to \infty$, and $\tau^* \leq \tau^R$. If $\lambda < 1$, the welfare optimum can be a subsidy. Appendix A.2 verifies the second order condition holds for $\lambda > \frac{1-s}{2}$.

2.2 What about the level of sensitivity to tax rates?

The foregoing discussion provides a baseline revenue maximizing rate of one half. But, this made a strong assumption about the sensitivity of investment to tax rates. That result immediately collapses if we introduce a degree of elasticity into the cost function, $C(e)$. Here, we incorporate the elasticity of investment effort to the keep rate (i.e., 1 – tax rate) Specifically, the foregoing discussion assumes quadratic cost $C(e) = \frac{k}{2}e^2$. As indicated above, the optimal effort in this case is $e^* = \beta A/k$ and the expected base $E[Y] = Ae^* = \left(\frac{A^2}{k}\right)\beta$ are linear in the retained share $\beta = 1 - \tau$. The elasticity of investment effort with respect to the net-of-tax rate is therefore identically one and the revenue maximizing rate is pinned at one half. However, this unit elasticity is a function of the quadratic cost function.

This section explores what happens when we replace quadratic cost with an iso-elastic (power) cost. Here, we set the cost function as:

$$C(e) = \frac{ke^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}}, \qquad \nu > 0 \tag{7}$$

The foregoing result arises if $\nu = 1$, in which case the cost function is the quadratic cost function. Now, in the power cost function, $\frac{dC(e)}{de} = ke^{\frac{1}{\nu}}$. Thus, the effort condition $\beta A = C'(e)$ yields:

$$e^*(\beta) = \left(\frac{\beta A}{k}\right)^{\nu}, \qquad E[Y] = Ae^* = A\left(\frac{A}{k}\right)^{\nu}\beta^{\nu}, \qquad \frac{d\ln E[Y]}{d\ln\beta} = \nu \tag{8}$$

The latter expression, $\frac{d\ln E[Y]}{d\ln\beta} = \nu$, is notable. This highlights that, in our context, ν represents the sensitivity of the return-base $E[Y]$ with respect to how much of the base the investor gets to keep (i.e., β). But, because the return base is simply Ae (i.e., investment-effort multiplied by the return parameter, which we take to be a constant), ν is the sensitivity of investment effort to the net-of-tax rate $(1-\tau)$. Intuitively, this is the responsiveness of investment-effort to how much the investor retains. This has significant implications for the optimal tax rates. Here, the parameter ν is the elasticity of the tax base with respect to the net-of-tax rate. The direct revenue is now $R(\tau) = \tau E[Y] = A\left(\frac{A}{k}\right)^{\nu}\tau(1-\tau)^{\nu}$. Differentiating with respect to τ yields an optimal revenue maximizing rate of:

$$\tau^R = \frac{1}{1+\nu}$$

This gives us a Laffer-style result, reminiscent of the elasticity documented in Saez (2001) in relation to income taxes: if investors respond more strongly to incentives, then the optimal tax rate is lower. Only if the elasticity of investment-effort to the keep rate is one (i.e., $\nu = 1$) would we get a tax on real capital gains of one half of *real* gains. This begs the question of what an appropriate value for ν is. Commonly, the Laffer peak is regarded as around 25-30% of nominal gains (generally ignoring the fiscal spillover, growth, and downstream impact of investment) This would imply $\nu \approx 2.3$ to 3. A Laffer peak of up to 40% (Agersnap and Zidar, 2021),[3] would imply $\nu = 1.5$. Alternatively, the model in Strulik and Trimborn (2012) suggests

[3] Agersnap and Zidar (2021) present a range of potential Laffer peaks, with the peak going up to 40%. Note that this peak is best thought of as the peak in a situation where there are no fiscal spillovers, or downstream

that the Laffer peak for capital gains tax is zero percent, which would imply $\nu \to \infty$. However, estimates vary significantly. The significant variation in empirical Laffer estimates, and the possibility that human and financial capital are increasingly mobile, would suggest erring towards a higher, rather than lower, estimate for ν.

The welfare maximizing rate also changes when we incorporate the elasticity of investment effort to the keep rate, ν. The certainty equivalent at the optimum is $CE^* = \frac{\frac{1}{1+\nu}(\beta A)^{1+\nu}}{k^\nu}$ (setting the risk premium to zero for exposition). Therefore, maximizing $W = CE^* + \lambda R$ over τ gives:

$$\tau^* = \frac{\lambda - 1}{\lambda(1+\nu) - 1} \tag{9}$$

Notably, this collapses to the baseline rate of $\frac{\lambda-1}{2\lambda-1}$ if $\nu = 1$ (i.e., we are in the quadratic cost case). Further, the rate is increasing in λ (i.e., the value of public revenue relative to private welfare). Alternatively, the welfare-maximizing rate tends to zero as $\lambda \to 1$. We capture these results in the following proposition.

Proposition 2 (Optimal rates under general base elasticity of investment-effort to the keep rate): Let the effort cost be the isoelastic form $C(e) = \frac{ke^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}}$ with $\nu > 0$ so that the expected base $E[Y] = Ae^*$ has constant elasticity ν with respect to the net of tax rate $\beta = 1 - \tau$. Then, (i) the revenue maximizing rate is $\tau^R = \frac{1}{1+\nu}$; (ii) the welfare maximizing rate (with $s = 0$) is $\tau^* = \frac{\lambda-1}{\lambda(1+\nu)-1}$; (iii) both rates are strictly decreasing in ν and both coincide with the baseline rates of Proposition 1 at $\nu = 1$. As $\nu \to \infty$ both rates tend to zero. As $\nu \to 0$, the revenue maximizing rate tends towards one.

---

effects, as their focus is on the elasticity of capital realizations to tax rates. However, we also acknowledge that these empirical Laffer peaks reflect realized behavior and the possibility of strategically timed gains and losses, which are not explicitly captured in the focal optimal rates.

We can also consider more closely the relationship between the optimal tax rate and the investment-effort sensitivity to the keep rate, ν. So, differentiating τ^R with respect to ν yields $\frac{d\tau^R}{d\nu} = -\frac{1}{(1+\nu)^2}$. Similarly, differentiating τ^* with respect to ν yields $\frac{d\tau^*}{d\nu} = \frac{-\lambda(\lambda-1)}{(\lambda+\lambda\nu-1)^2}$. Notably both are negative: the higher the elasticity the lower the revenue-maximizing and welfare-maximizing rates. This has some interesting implications if we consider different investor types. The core analysis homogenizes investors. However, the empirical evidence shows that superstar inventors (Akcigit et al., 2016), and superstar athletes (Kleven et al., 2013), are especially likely to move overseas, implying a higher sensitivity to tax rates. Our model helps to theoretically ground this insight: the more sensitive the individual is to tax rates, the lower is the required tax rate to drive investment-effort.

2.3 What about the treatment of losses?

The foregoing discussion assumes losses are treated symmetrically. The investor keeps $(1-\tau)Y$ of any realized gain Y, and a loss of $Y < 0$ generates a tax refund equal to $\tau|Y|$. This is akin to a full-loss-offset per Domar and Musgrave (1944). This might be defensible if we assume investors run operating companies and/or can fully offset losses, such as might be the case if the investor has a total net gain against which to offset the losses. However, real tax systems typically have a partial offset. For example, a loss accrued in a past period is not indexed to inflation, meaning that its real value decreases over time.

We explore the impact of such a partial loss offset by introducing a loss-offset ratio $\ell \in [0,1]$. Here, the government refunds losses at a rate of $\ell\tau$. Thus, if $\ell = 1$, we have the symmetric baseline and if $\ell = 0$ investors cannot offset losses. Therefore, the investor bears more downside risk. This manifests in the investor's risk retention rising from $(1-\tau)$ to $(1-\ell\tau)$.

Therefore, assuming quadratic cost for expositional simplicity, the certainty equivalent becomes:

$$CE(e) = \underbrace{(1-\tau)Ae}_{\text{Exp. retained gain}} - \overbrace{\frac{k}{2}}^{\text{Effort cost}}e^2 - \underbrace{\frac{a}{2}(1-\ell\tau)^2\sigma^2}_{\text{Risk premium}} \tag{10}$$

We note that the change in the risk premium is independent of investment-effort. Thus, the different loss treatment does not directly impact investment-effort in this baseline model. That is, the optimal effort remains as $e^* = \frac{(1-\tau)A}{k}$ and the base remains as $E[Y] = Ae^*$. The revenue-maximizing tax rate, τ^R, remains unchanged. In the quadratic effort-cost case, it stays at one half. Loss asymmetry operates primarily through whether the investor participates. To see this the certainty equivalent at the optimum, with $m = \frac{A^2}{k}$ and $s = \frac{a\sigma^2}{m}$ is:

$$CE^* = \frac{m}{2}[(1-\tau)^2 - s(1-\ell\tau)^2] \tag{11}$$

This is strictly decreasing in loss asymmetry (i.e., increasing in ℓ) since $\frac{dCE^*}{d\ell} = a\sigma^2\tau(1-\ell\tau) > 0$. That is, the more of the loss the investor can offset, the higher is the certainty equivalent. Worse loss offsets lower the surplus, which lowers the participation threshold. The investor therefore participates only if $CE^* \geq \bar{u}$. That is, we require:

$$(1-s\ell^2)\tau^2 - 2(1-s\ell)\tau + \left(1 - s - \frac{2\,\bar{u}}{m}\right) \geq 0$$

We can solve to obtain $\tau_{\text{participation}}(\ell)$ as the rate below which investors participate in investing. The intuition is that as the offset worsens (i.e., ℓ decreases), investors exit at lower rates. The implication for the optimal tax rate runs through the exit channel. Assuming a quadratic cost function, the effective revenue maximizing rate is thus $\min\left\{\frac{1}{2}, \tau_{\text{participation}}(\ell)\right\}$. We capture this in Figure 2, which plots the relationship between the participation tax rate and the proportion of losses that can be offset. The figure assumes the quadratic cost function (on which, see caveats, above). The figure highlights that the greater the portion of losses that can be offset, the greater the potential capital gains tax that can be supported. We also note this effect in the following proposition:

Proposition 3 (Loss offset and participation): Let the government refund losses at a rate $\ell\tau$, $\ell \in [0,1]$. Then $\ell = 1$ represents full symmetric loss offsetting. Then, the investor's effective risk retention is $1 - \ell\tau$ and $CE(e) = (1-\tau)Ae - \frac{k}{2}e^2 - \frac{a}{2}(1-\ell\tau)^2\sigma^2$. Then, (i) optimal effort is $e^* = \frac{(1-\tau)A}{k}$ and the expected base are independent of ℓ, so that the interior revenue-maximizing rate remains at one half in a quadratic cost situation; (ii) the certainty equivalent $CE^* = \frac{m}{2}[(1-\tau)^2 - s(1-\ell\tau)^2]$ is strictly decreasing in loss asymmetry, so the participation threshold $\bar{\tau}_{\text{participation}}(\ell)$ is strictly increasing in the proportion of losses that the investor can offset, ℓ; (iii) the effective revenue-maximizing rate is $\min\left\{\frac{1}{2}, \bar{\tau}_{\text{participation}}(\ell)\right\}$, which falls below one half once offset is sufficiently incomplete.

Corollary 3.1 (Losses and exit): Imperfect loss offset lowers the tax rate at which the investor exits. It therefore reinforces the mobility/participation channel. Poor loss treatment, like a lightly taxed offshore alternative, drives investors out and caps the feasible tax rate below one half.

[Insert Figure 2 about here]

3 What about indirect benefits to capital: fiscal spillovers

The initial baseline analysis is myopic: it ignores the indirect benefits of capital investment. For example, capital can help to drive corporate activity and employment. This is the case whether the investing is in the form of human capital (i.e., as a founder), primary shares (i.e., as a venture capitalist), secondary shares (i.e., as an investor in the stock market), or in property. Furthermore, capital gains tax rates can influence firms' cost of capital (Guenther and Willenborg, 1999; Huizinga et al., 2018), which in turn influences their ability to raise capital to invest. For present purposes, we do not distinguish between investment types. Rather, it suffices that such investment generates growth, which results in ancillary revenue from corporate income tax, payroll tax, employment, etc.

We separate the growth-benefits into two types: (a) direct benefits associated with the investment per se, and (b) broader societal growth benefits. This section deals with the direct benefits of investing. We address broader long-term growth impacts in Section 5.

We capture this by letting each unit of expected activity, $E[Y]$ generate ancillary revenue, $\phi E[Y]$. For expositional convenience, we use a quadratic cost function, but note the above limits of such a function. We can also generalize this over future time periods, such that for all future time periods t, the ancillary revenue is:

$$\phi = \sum_{t \geq 1} \omega^t \psi = \frac{\psi}{\rho}, \qquad \omega = \frac{1}{1+\rho}$$

Where, ψ is the per period ancillary yield and ρ is the discount rate. We treat ϕ as a fiscal externality. The investor responds only to their own retention (i.e., to τ). Thus, the optimal effort e^* is unchanged. However, the government now collects $G(\tau) = (\tau + \phi)E[Y] = (\tau + \phi)(1 - \tau)m$, where $m = \frac{A^2}{k}$. Intuitively, this changes the government's optimization solution as the government must now consider tax revenue from future growth.

We can differentiate the government's revenue with respect to τ to obtain

$$G'(\tau) \propto (1-\tau) - (\tau+\phi) = 1 - 2\tau - \phi$$

Taking the first order condition yields:

$$\tau_{\phi}^{R} = \frac{1-\phi}{2} \leq \frac{1}{2} \tag{12}$$

This tells us that if we consider additional tax revenue from growth, then the 'revenue maximizing' tax rate is always before the rate that only considers first order, or static, revenue effects. Only if we consider there to be no growth-related externalities would the tax rate tend towards 50%. But, assuming no growth related externalities is unrealistic and contrary to the overwhelming weight of empirical evidence. Even at an intuitive level, startups, which rely on investors for funding, generate innovation, which generates economic growth, which generates additional tax revenue.

The welfare optimal rate is now lower as well. Recall that the welfare is $W = CE^* + \lambda G$, where λ represents the relative weight given to private versus public capital. Let us simplify and assume that $s = \frac{a\sigma^2}{m} = 0$, for expositional clarity. We then see that:

$$\frac{W'(\tau)}{m} = -(1-\tau) + \lambda[(1-\tau) - (\tau+\phi)]$$

Taking the first order condition and solving yields:

$$\tau_{\phi}^{*} = \frac{\lambda - 1 - \lambda\phi}{2\lambda - 1} \tag{13}$$

But, we can rearrange this in terms of the foregoing naïve welfare maximizing rate of $\tau^*_{\text{Naive, Welfare}} = \frac{\lambda-1}{2\lambda-1}$. This means that:

$$\tau_\phi^* = \frac{\lambda - 1 - \lambda\phi}{2\lambda - 1} = \tau^*_{\text{Naive, Welfare}} - \frac{\lambda\phi}{2\lambda - 1}$$

We capture this in the following proposition:

Proposition 4 (Optimal rates with a fiscal spillover). Let each unit of expected gain generate fiscal spillover $\phi \geq 0$, so government revenue is $G(\tau) = (\tau + \phi)(1 - \tau)m$, where $m = \frac{A^2}{k}$. Then with $s = 0$, (i) $\tau_\phi^R = \frac{(1-\phi)}{2}$, (ii) $\tau_\phi^* = \frac{\lambda-1-\lambda\phi}{2\lambda-1} = \tau^*_{\text{Naive, Welfare}} - \frac{\lambda\phi}{2\lambda-1}$. Both rates are strictly decreasing in ϕ and weakly below their Proposition 1 counterparts.

The intuition is clear: investment-driven growth lowers the optimal CGT rate. Taxing more suppresses investment effort and deployment. Such a tax reduces not just the straightforward CGT base but the ancillary base. Indeed, in this framework, the government might *subsidize* investment (i.e., via tax credits). To find this, we solve for the level of ϕ at which the tax rate would cross into becoming negative. This occurs at:

$$\phi^\dagger = \frac{\lambda - 1}{\lambda}$$

Corollary 4.1 (Subsidy threshold). The welfare-maximizing policy is a subsidy ($\tau_\phi^* < 0$) if and only if $\phi > \phi^\dagger = (\lambda - 1)/\lambda$. When the fiscal spillover exceeds the share of social value the government places on private surplus, optimal policy credits rather than taxes investment.

The intuition is that if employment and economic activity are sufficiently high, then the policy is to encourage investment in such activities rather than tax them. This is reminiscent of the case for R&D tax incentives (see e.g., Chen and Hshieh, 2024), and venture tax incentives broadly (Campello and Junqueira, 2025; Keuschnigg and Nielsen, 2004). This reflects the logic behind various startup-funding incentives. Examples include Australia's Early Stage

Venture Capital Limited Partnership (ESVCLP) and Early Stage Innovation Company (ESIC) programs: these both involve tax credits for investing in qualifying early stage companies.

We can also consider the situation under isoelastic cost (as opposed to quadratic cost). So, to do this, we define the cost function as $C(e) = \frac{ke^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}}$. We carry the elasticity ν throughout the fiscal spillover analysis. Under this isoelastic effort-cost function, the investor's optimal investment-effort is $e^* = \left(\frac{\beta A}{k}\right)^{\nu}$ and the expected base is $E[Y] = A\left(\frac{A}{k}\right)^{\nu}\beta^{\nu}$, as in Section 2.2. Then, writing $m_\nu = A\left(\frac{A}{k}\right)^{\nu}$, so that $m_1 = A\left(\frac{A}{k}\right)$ recovers the quadratic case, the government collects:

$$G(\tau) = (\tau + \phi)m_\nu(1-\tau)^{\nu}$$

Differentiating G with respect to τ and setting to zero gives $(1-\tau) = \nu(\tau+\phi)$, so the revenue maximizing rate is:

$$\tau_\phi^R = \frac{1-\nu\phi}{1+\nu}$$

Maximizing $W = CE^* + \lambda G$ (and setting $s = 0$ for clarity) yields the welfare maximizing rate of:

$$\tau_\phi^* = \frac{(\lambda-1) - \lambda\nu\phi}{\lambda(1+\nu)-1}$$

Both rates collapse to the quadratic case when $\nu = 1$ and collapse to the baseline case when $\phi = 0$. The economic content of the generalization is that the elasticity of investment effort to the keep rate scales how hard the spillover bites. Without a fiscal spillover, the revenue-maximizing rate remains at $\frac{1}{1+\nu}$ rather than one half. A more responsive base (i.e., a larger ν) amplifies the impact of the spillover, reducing the welfare-maximizing rate. Furthermore, τ_{ϕ}^{*} turns negative at the generalized subsidy threshold of:

$$\phi^{\dagger} = \frac{\lambda - 1}{\lambda \nu}$$

Corollary 4.2 (Spillover rate under general base keep rate elasticity): Under the isoelastic cost structure, the spillover-adjusted revenue-maximizing rate is $\tau_{\phi}^{R} = \frac{1-\nu\phi}{1+\nu}$ and the welfare-maximizing rate is $\tau_{\phi}^{*} = \frac{(\lambda-1)-\lambda\nu\phi}{\lambda(1+\nu)-1}$. Both are strictly decreasing in the elasticity, ν, of investment effort to the keep rate. The welfare-maximizing rate becomes a subsidy, and $\tau_{\phi}^{*} < 0$ at $\phi^{\dagger} = \frac{\lambda-1}{\lambda\nu}$. The subsidy threshold is decreasing in ν, implying that a higher elasticity reduces the level of spillover necessary for a subsidy to be deployed. The foregoing quadratic results are the special case when $\nu = 1$.

4 What about differently taxed, or untaxed, alternatives?

The foregoing analysis assumes that the investor has one place to exert effort or deploy capital. It also implicitly assumes that the investor has a binary choice: stay and pay the CGT or move to an untaxed jurisdiction. This is operationalized via the participation constraint. However, investors, corporations, and individuals can, and do, move between locations. For example, in the corporate field, corporate taxes significantly influence corporate HQ relocations (Chow et al., 2022). Holmstrom and Milgrom (1991) provide a framework through which to analyze what happens when the investor faces multiple different tax rates. This could arise if different assets are taxed at different rates or if they can access different jurisdictions with different rates. This is also akin to the notion of an investor moving to an untaxed, which could be considered a separate 'task' that the investor might pursue.

We can operationalize this much like in the multi-task setting in Holmstrom and Milgrom (1991). Here, we assume that there are J different places for the capital (think different assets with different tax rates or different jurisdictions). Each channel i produces a payoff $Y_i = A_i e_i + \epsilon_i$. Each investment has a potentially different tax rate, with the investor retaining $\beta_i = 1 - \tau_i$ of the gains from the investment. We assume that effort is jointly costly. In this section, for tractability, we use a quadratic cost structure. We also assume that the returns are potentially correlated. This gives us the following:

$$C(e) = \frac{1}{2}\boldsymbol{e}^T\boldsymbol{K}\boldsymbol{e}, \qquad \boldsymbol{K} = \begin{pmatrix} k_1 & \kappa \\ \kappa & k_2 \end{pmatrix}$$

Where there are two channels, k_i denotes the marginal cost of effort for the i^{th} investment, and the off-diagonal κ is the interaction. Thus, $\kappa > 0$ means that the channels are substitutes. For example, investing more in one asset raises the marginal cost of the other, for example due to limited attention issues. This is akin to the issues documented within VC/PE, such as diseconomies of scale (Bernile et al., 2007; Cumming and Dai, 2011; Humphery-Jenner, 2012; Lopez-de-Silanes et al., 2015), and the appropriate level of diversification (Humphery-Jenner, 2013; Knill, 2009). Alternatively, if $\kappa < 0$ the different investment channels are complements. This is intuitively akin to diversification or knowledge sharing across asset classes.

The investor therefore maximizes $\sum_i \beta_i A_i e_i - C(e)$. The first order condition is given by $\boldsymbol{K}\boldsymbol{e}^* = \boldsymbol{b}$, with $\boldsymbol{b} = \left(\beta_1 A_1, \dots, \beta_J A_J\right)^T$. Therefore, $\boldsymbol{e}^* = \boldsymbol{K}^{-1}\boldsymbol{b}$.

For the sake of illustration, let us assume that the investment channels are symmetric with $k_1 = k_2 = k$ and $A_i = A$. In this scenario, the only difference between the assets is the amount of tax faced. This is akin to a situation where an investor can seamlessly shift assets into a low taxed jurisdiction. Solving $\boldsymbol{e}^* = \boldsymbol{K}^{-1}\boldsymbol{b}$ gives the following:

$$e_1^* = \frac{A(k\beta_1 - \kappa\beta_2)}{k^2 - \kappa^2}, \qquad e_2^* = \frac{A(k\beta_2 - \kappa\beta_1)}{k^2 - \kappa^2} \tag{14}$$

This yields an intuitive result. A higher tax on asset one leads the investor to shift towards asset two. To see this, consider the derivative of e_1^* with respect to β_2. This captures the amount of effort/investment devoted to asset one when the investor retains more of asset two.

$$\frac{de_1^*}{d\beta_2} = -\frac{A\kappa}{k^2 - \kappa^2} < 0$$

The intuition behind this is that taxing channel two less (i.e., by increasing β_2) results in effort being diverted away from channel 1.

Lemma 3 (Multitask effort allocation). With J channels, joint cost $C(e) = \frac{1}{2}\boldsymbol{e}^T\boldsymbol{K}\boldsymbol{e}$, with $\boldsymbol{K}$ being symmetric positive definite, and retentions $\beta_i = 1 - \tau_i$, optimal effort is $\boldsymbol{e}^* = \boldsymbol{K}^{-1}\boldsymbol{b}$. In the symmetric two channel case ($k_i = k, A_i = A$, off-diagonals are κ, and the substitution index is $\rho_s = \frac{\kappa}{k} \in [0,1)$), $e_1^* = \frac{A(k\beta_1-\kappa\beta_2)}{k^2-\kappa^2}$, $\frac{de_1^*}{d\beta_2} = -\frac{A\kappa}{k^2-\kappa^2} < 0$, so effort is diverted toward the more lightly taxed channel. Positive definiteness conditions are in Appendix A.3.

4.1 Untaxed alternatives

We can now explore what happens when one channel has zero tax. For simplicity, we will initially ignore the aforementioned 'fiscal spillover' or growth impact of investment. In our two channel case, we let $\tau_2 = 0$; and thus, $\beta_2 = 1$. Examples of this depend on the jurisdiction. For example, in Australia, there is no capital gains tax on investments made within an ESVCLP or in an ESIC-qualifying company. Furthermore, there is no CGT on an individual's principal place of residence. Another example would be for individuals who can shift assets into a zero tax jurisdiction, as is allowed in some countries.

The government faces the following situation. The revenue raised is simply a function of the taxed asset, not the untaxed asset. Thus, the direct revenue is $R = R(\tau_1) = \tau_1 A e_1^*$. Welfare is $W = CE^* + \lambda R$, with $CE^* = \frac{1}{2}\sum_i \beta_i A_i e_i^*$. Now, let us define the 'substitution index', for notational convenience, as $\rho_s = \frac{\kappa}{k} \in [0,1)$. This captures the extent to which the assets are substitutes. We can now obtain the relevant tax rates.

The initial revenue maximizing rate we obtain by substituting e_1^* into $R(\tau_1) = \tau_1 A e_1^*$, differentiating with respect to τ_1 and setting equal to zero. This gives us:

$$\tau_1^R = \frac{1}{2}(1 - \rho_s) \tag{15}$$

If we consider the 'fiscal spillover', or growth impact, of investment, then the government's return to investment becomes $R(\tau_1) = (\tau_1 + \phi) A e_1^*$. In this case, the revenue-maximizing tax rate would be:

$$\tau_1^R = \frac{1}{2}(1 - \rho_s) - \frac{\phi}{2} \tag{16}$$

The welfare maximizing rate derives from differentiating $W = CE^* + \lambda R$ with respect to τ_1 and setting to zero. This yields:

$$\tau_1^* = (1 - \rho_s)\frac{\lambda - 1}{2\lambda - 1} \tag{17}$$

Similarly, if we consider the growth-impact of investment, or fiscal spillover, then welfare is $W = CE^* + \lambda G$, where $G = (\tau_1 + \phi)Ae_1^*$. Differentiating W with respect to τ_1 and setting equal to zero then yields:

$$\tau_1^* = (1-\rho_s)\frac{\lambda-1}{2\lambda-1} - \frac{\lambda\phi}{2\lambda-1} \tag{18}$$

The foregoing has several clear implications. As $\rho_s \to 1$, the assets become perfect substitutes. The investor can seamlessly shift between the taxed and untaxed good. In this case, the revenue-maximizing tax rate approaches zero. Further, as $\rho_s \to 1$, the welfare-maximizing tax rate especially approaches zero, particularly when we consider fiscal spillovers, or growth-effects, of investment.

Proposition 5 (Optimal rate with a lower-taxed alternative). With a second, untaxed channel ($\tau_2 = 0$) and substitution index $\rho_s \in [0,1)$: (i) revenue-maximizing rate $\tau_1^R = \frac{1-\rho_s}{2}$, and with a fiscal spillover of ϕ, is $\frac{1-\rho_s-\phi}{2}$; (ii) welfare maximizing rate is $\tau_1^* = \frac{(1-\rho_s)(\lambda-1)}{2\lambda-1}$, or with a fiscal spillover of ϕ is $\tau_1^* = (1-\rho_s)\frac{\lambda-1}{2\lambda-1} - \frac{\lambda\phi}{2\lambda-1}$. All rates are strictly decreasing in ρ_s.

Corollary 5.1 (Mobility limit). As $\rho_s \to 1$ (the channels become perfect substitutes; frictionless mobility or a fully fungible untaxed asset), both the revenue- and welfare-maximizing rates converge to zero. With fiscal spillovers, the unconstrained optimum can be negative.

The substitution result has further implications for investor mobility. In Section 2.2, we considered how investor skill can manifest in the sensitivity to the keep rate. Here, we consider an alternative avenue through which taxes can cause the tax base to bleed superstar investors. Specifically, given that the empirical literature shows that skilled people are more likely to move, this suggests that skilled individuals face a higher ρ_s term, suggesting a greater substitutability between exerting effort domestically or deploying their skills overseas. This might arise because the individual's skillset makes the low taxed overseas option more attractive, increasing its substitutability. Thus, our model helps to illustrate a channel through which countries lose skilled investors.

4.2 What if opportunities are taxed equally?

If all investment alternatives are taxed uniformly, then there is no intuitive reason for investors to substitute between similar assets. To see this, set $\tau_1 = \tau_2 = \tau$. In this case $e_1^* = e_2^* = \frac{A\beta}{k+\kappa}$. Total revenue is then $2\tau A e^* = \frac{2A^2}{k+\kappa}\tau(1-\tau)$. In this case, ignoring fiscal spillover, or growth-effects of investment, we obtain the same naïve revenue maximizing tax rate of 50%.

The intuition behind this result manifests in why governments might apply the same uniform capital gains tax to all asset classes. Were they to not do so, then for assets with similar characteristics, capital would flow to the lower taxed alternative. This also underscores why some jurisdictions attempt to tax capital held within overseas organizations.

Corollary 5.2 (Uniform-taxation neutrality). If all channels are taxed at a common rate $(\tau_1 = \tau_2 = \tau)$, substitution is neutralized: $e^* = \frac{A\beta}{k+\kappa}$ and the revenue-maximizing rate returns to 1/2 (absent spillovers). Uniform taxation of like assets is the policy that forecloses base-shifting.

4.3 What about consuming now versus investing?

We can now explicitly analyze what happens when the individual has a channel for their money that is an alternative to investing: consuming now. Consuming now has no effort disutility and yields utility that is worth a fraction $\delta \in (0,1]$ of investing. This is the now-versus-later trade-off. Such consumption is taxed at a separate rate, τ_c, representing sales tax and the corporate tax footprint of that spending. The investor's certainty equivalent then becomes:

$$CE(e,c) = \underbrace{\beta_k Ae - \frac{k}{2}e^2}_{\text{Invest}} + \delta \underbrace{\left[\beta_c Bc - \frac{k_c}{2}c^2\right]}_{\text{Consume}} - \underbrace{\kappa ec}_{\text{Substitution}} \qquad (19)$$

Where, $\beta_k = 1 - \tau_k$, $\beta_c = 1 - \tau_c$. Further, the term $\frac{k_c}{2}c^2$ denotes the diminishing marginal utility of current consumption (not a cost). Therefore, the δ multiplier captures the value of consumption relative to investing. Intuitively, a low δ means that the investor prefers investing whereas a high δ means that the investor prefers consumption. In the above, the term κec represents resource competition; the activities (i.e., investing and consuming) are substitutes for $\kappa > 0$.

The first order conditions are therefore:

$$e: \beta_k A = ke + \kappa c, \qquad c: \delta\beta_c B = \delta k_c c + \kappa e$$

This leads to the following solutions:

$$e^* = \frac{\delta(k_c\beta_k A - \kappa\beta_c B)}{\delta k k_c - \kappa^2}, \qquad c^* = \frac{k\delta\beta_c B - \kappa\beta_k A}{\delta k k_c - \kappa^2}$$

The result is clear: a higher CGT rate τ_k results in less investment effort and results in more present consumption.

4.3.1 Optimal CGT rates, ignoring the growth impact of investing

We can now analyze the optimal CGT rate. We start by ignoring the 'fiscal spillover', or growth effect, of investment. Here, government revenue derives from the CGT on investment gains and the sales tax on consumption. Thus, government revenue is $R = \tau_k Ae^* + \tau_c Bc^*$. In the symmetric case where $B = A$ and $k_c = k$ (with $\rho_s = \kappa/k$), maximizing R over τ_k gives the revenue maximizing rate, and maximizing over $W = CE^* + \lambda R$ the welfare maximizing rate. This yields the following 'naïve' revenue maximizing rate (ignoring the growth impact of investing):

$$\tau_k^R = \frac{1}{2}\left[1 - \rho_s\left((1 - \tau_c) - \frac{\tau_c}{\delta}\right)\right] \tag{20}$$

Similarly, we obtain the following welfare maximizing rate:

$$\tau_k^* = \frac{(\lambda - 1)[1 - \rho_s(1 - \tau_c)] + \frac{\rho_s \lambda \tau_c}{\delta}}{2\lambda - 1} \tag{21}$$

The foregoing has a counter-intuitive result: the higher the consumption tax, τ_c, the higher the potential CGT. The intuition is that if the government raises consumption tax by ϵ, then the government could also raise CGT by an equivalent amount while maintaining preferences between investing and consumption. But, this raises the clear issue: investing can generate economic growth, which can drive future tax revenue. We capture this effect in the following section.

4.3.2 Optimal CGT including the fiscal spillover

We can now consider a situation where the investment contains downstream tax benefits, ϕ (i.e., employment, corporate tax, etc.). The government revenue then becomes $R = (\tau_k + \phi)Ae^* + \tau_c Bc^*$. From here, we can identify the government-revenue maximizing tax rates and the welfare maximizing rates. The government revenue maximizing rate becomes:

$$\tau_k^R = \frac{1}{2}\left[1 - \phi - \rho_s\left((1 - \tau_c) - \frac{\tau_c}{\delta}\right)\right] \tag{22}$$

The intuition is similar to in the foregoing sections: the revenue maximizing CGT rate declines when such capital investment generates tax revenue indirectly. The welfare maximizing rate similarly falls. The intuition is similar: if capital investment generates additional fiscal spillovers, which increase both private surplus and government revenue, then it is optimal to tax capital at a lower rate than would otherwise be the case. Thus, the welfare maximizing rate becomes:

$$\tau_k^* = \frac{(\lambda - 1)[1 - \rho_s(1 - \tau_c)] + \frac{\rho_s \lambda \tau_c}{\delta} - \lambda\phi}{2\lambda - 1} \tag{23}$$

Proposition 6 (Optimal CGT with a consumption margin). Let investing compete with present consumption, valued at $\delta \in (0,1]$ and taxed at τ_c, with symmetric primitives $B = A$, $k_c = k$, and $D = \delta k^2 - \kappa^2 > 0$. Equivalently $\rho_s^2 < \delta$ (i.e., $\rho_s < \sqrt{\delta}$). Then, (i) revenue maximizing rate $\tau_k^R = \frac{1}{2}\left[1 - \phi - \rho_s\left((1 - \tau_c) - \frac{\tau_c}{\delta}\right)\right]$; (ii) welfare-maximizing rate $\tau_k^* = \frac{(\lambda-1)[1-\rho_s(1-\tau_c)]+\frac{\rho_s\lambda\tau_c}{\delta}-\lambda\phi}{2\lambda-1}$.

Corollary 6.1 (Consumption-tax complementarity). Both rates are strictly increasing in the consumption-tax rate τ_c: $\frac{d\tau_k^R}{d\tau_c} = \frac{1}{2}\rho_s\left(1 + \frac{1}{\delta}\right) > 0$ and $\frac{d\tau_k^*}{d\tau_c} = \frac{\rho_s\left[(\lambda-1)+\frac{\lambda}{\delta}\right]}{2\lambda-1} > 0$. A higher tax on present consumption raises the optimal capital gains tax because it preserves the relative price between consuming now and investing.

5 Optimal CGT when investment drives growth

The foregoing discussion ignores the impact of investment on growth. Ignoring such growth impacts causes the model to understate the cost of taxing investment. To analyze this, we credit each expected investment gain with the present value of future tax revenue that the investment sets in motion. This is consistent with the well-documented impact of investment on growth. For example, Borensztein et al (1998) find that foreign direct investment drives economic growth, including through technology transfer. Xu (2023) and Lindsey and Stein (2026) demonstrate a clear link between angel investment and business formation, which also increases employment. Similarly, Arefeva et al (2025) find a link between capital taxation and job growth, showing a clear link between capital tax cuts and labor growth. Further, even secondary asset markets (i.e., the secondary market for shares or existing property) is inextricably linked to prices in the primary market (Mauer and Senbet, 1992), which themselves influence access to capital and growth.

This section aims to explicitly incorporate the growth-impact of investing. To be clear, in the foregoing discussion of fiscal spillovers, ϕ is best thought of as the contemporaneous ancillary tax on *this* investment's activity, whereas g is the tax on the *additional* compounding output it sets in motion.

In our framework, if a unit of deployed capital generates output at return ב, taxed across all bases at a blended effective rate of θ, with the reinvested base growing at a rate γ, discounted at a rate $i > \gamma$, then we have the following growth dividend that accrues to the government.

$$g = \sum_{t\geq 1} \frac{(1+\gamma)^{t-1}\theta ב}{(1+i)^t} = \frac{\theta ב}{i-\gamma} \tag{24}$$

The forgoing growth impact is effectively the perpetuity formula applied to future growth (and related tax) generated by the capital investment. As is standard for such a formula, as the growth rate (γ) approaches the discount rate (i), the total value of growth explodes.

We can then identify the government-revenue maximizing tax rate (τ_k^R), and the welfare maximizing tax rate (τ_k^*) as:

$$\tau_k^R = \frac{1}{2}\Big[1 - g - \rho_s((1-\tau_c) - \frac{\tau_c}{\delta})\Big] \tag{25}$$

$$\tau_k^* = \frac{(\lambda - 1)[1 - \rho_s(1-\tau_c)] + \frac{\rho_s \lambda \tau_c}{\delta}}{2\lambda - 1} - \frac{\lambda g}{2\lambda - 1} \tag{26}$$

Proof: To get the optimal tax rates, write $\beta_k = 1 - \tau_k$ and $D = \delta k^2 - \kappa^2$. This is for notational convenience. In the symmetric case (i.e., $A = B$ and $k_c = k$), we get:

$$Ae^* = \frac{\delta A^2}{D}(k\beta_k - \kappa\beta_c), \qquad Ac^* = \frac{A^2}{D}(\delta k\beta_c - \kappa\beta_k)$$

To obtain the revenue-maximizing rate: With the growth dividend, the revenue is $R = (\tau_k + g)Ae^* + \tau_c Ac^*$*. Multiplying through by the positive constant* $\frac{D}{A^2}$*, we obtain:*

$$\frac{D}{A^2}R = (\tau_k + g)\delta(k\beta_k - \kappa\beta_c) + \tau_c(\delta k\beta_c - \kappa\beta_k)$$

We then differentiate with respect to τ_k *and recall that* $\frac{d\beta_k}{d\tau_k} = -1$ *while* β_c *is fixed. Thus, we obtain* $\tau_k^R = \frac{1}{2}\left[1 - g - \rho_s((1 - \tau_c) - \frac{\tau_c}{\delta})\right]$.

For the welfare-maximizing rate, we maximize $W = CE^* + \lambda R$*, where* λ *represents the relative social value of a public dollar. Thus, we obtain*

$$\frac{dW}{d\tau_k} = (\lambda - 1)Ae^* - \lambda(\tau_k + g)\frac{\delta k A^2}{D} + \lambda\tau_c\frac{\kappa A^2}{D}$$

Substituting $Ae^* = \frac{\delta A^2}{D}(k\beta_k - \kappa\beta_c)$ *and multiplying by* $\frac{D}{A^2}$*, we obtain:*

$$(\lambda - 1)\delta(k\beta_k - \kappa\beta_c) - \lambda(\tau_k + g)\delta k + \lambda\kappa\tau_c = 0$$

We can then rearrange and collect terms to obtain $\tau_k^* = \frac{(\lambda-1)[1-\rho_s(1-\tau_c)] + \frac{\rho_s\lambda\tau_c}{\delta}}{2\lambda-1} - \frac{\lambda g}{2\lambda-1}$.

A corollary of the foregoing discussion is that there is a tradeoff between the gains to consumption tax and growth. Specifically, driving investment to consumption (i.e., via a high capital tax rate) can generate consumption tax. However, it risks reducing future tax revenue by reducing the growth dividend.

Proposition 7 (Optimal CGT with the growth dividend). Let investment generate a growth dividend $g = \sum_{t\geq 1} \frac{(1+\gamma)^{t-1}\theta\beth}{(1+i)^t} = \frac{\theta\beth}{i-\gamma}$ with $i > \gamma$, the present value of the future tax stream it sets in motion. Then g enters the optimal rates exactly as ϕ does: $\tau_k^R = \frac{1}{2}\left[1 - g - \rho_s((1-\tau_c) - \frac{\tau_c}{\delta})\right]$, and $\tau_k^* = \frac{(\lambda-1)[1-\rho_s(1-\tau_c)] + \frac{\rho_s\lambda\tau_c}{\delta}}{2\lambda-1} - \frac{\lambda g}{2\lambda-1}$.

6 What about the lock-in effect: A productivity tax?

The foregoing analysis sidesteps the question of whether CGT disincentivizes investors from reallocating capital. But, higher CGT rates do influence equity realizations (Ehling et al., 2026; Fischer and Gallmeyer, 2016). Indeed, Jin (2006) identifies significant differences in selling behavior for funds with tax-sensitive clients as opposed to those with tax-exempt clients. This is important because capital allocation is an important function of the market. But, because the investor pays tax only when the gain is realized, investors might cling to an inferior asset if it takes “too long” to regain the tax paid on sale. We can consider this within a two-period setting. As we indicate below, high capital gains taxes can become akin to a ‘productivity tax’, as alluded to in Holden (2026); the high CGT can deter people from moving from less productive to more productive assets.

To consider the lock-in effect, we contemplate a sequential environment. The investor currently holds a low growth asset. They can sell it and reinvest in the high growth asset. Alternatively, they can hold the low growth asset to maturity. We consider two situations: (a) where the investor will want to consume at the end (so they sell whichever asset they have at the end of period two), and (b) where the investor never sells (i.e., because their family inherits the asset, they gain utility from wealth per se, or they access liquidity by borrowing against the asset.

To ground this, consider the following set-up: An investor who holds an asset with the current value P, cost basis of b, so the embedded unrealized gain is $G = P - b$. Or, in percentage terms $\theta_G = \frac{G}{P} \in [0,1)$. If the investor retains the asset, it grows at a low rate r_L. A superior asset is available, which would grow at $r_H > r_L$. The investor chooses to either hold the low growth asset or to switch to the higher growth asset. If the investor switches assets, the investor must pay capital gains tax at a rate of τ, for a total tax payment of τG. They reinvest the remainder (i.e., $P - \tau G$) in the new high growth asset.

6.1 Baseline scenario: the investor consumes at the end

We start by considering the situation where the investor wants to consume at the end of period two. For present purposes, we assume that the investor must sell the asset in order to consume and they consume the total value of their investment. If the investor holds the low growth asset, then it grows to $P(1 + r_L)$. The total capital gain is then $P(1 + r_L) - b$. Thus, if they sell at the end, then the investor pays tax of $\tau[P(1 + r_L) - b]$. Therefore, the investor's total wealth if they hold the low growth asset is:

$$W_{\text{hold}} = (1 - \tau)P(1 + r_L) + \tau b \tag{27}$$

The investor may alternatively sell the low growth asset and reinvest the proceeds. If the investor does this, they can reinvest $P - \tau G$, which then grows at r_H. The investor must then pay capital gains tax on the gain, resulting in a tax charge of $\tau r_H(P - \tau G)$. Therefore, if the investor switches assets, their total wealth is:

$$W_{\text{switch}} = [P - \tau G][1 + (1 - \tau)r_H] \tag{28}$$

Therefore, we can find the growth rate r_H at which the investor would switch by setting $W_{\text{hold}} = W_{\text{switch}}$. This yields:

$$r_H^* = \frac{r_L}{1 - \tau\theta_G} \tag{29}$$

The investor reallocates capital only if $r_H > r_H^*$. This means that there is set of assets for which $r_H^* > r_H > r_L$ for which switching is not rational. Long held, highly appreciated, assets are the most locked in. Furthermore, we can observe that the higher is the tax rate τ, the higher is the hurdle rate r_H^*

The consequence of the lock-in effect is to reduce the growth dividend discussed above. To see this, suppose that we have a population of investors, each of which have an embedded-gain fraction of θ_G. They differ in relation to their future opportunity r_H. We define the cumulative distribution function of r_H across investors as $F(x)$: the fraction of the population whose available opportunity is less than or equal to x. Therefore, $F(r_H^*)$ is the fraction of the population whose opportunity set is below the hurdle rate r_H^* and $1 - F(r_H^*)$ is the fraction above the hurdle rate: this is the proportion of people who switch assets.

We then define the aggregated realized gain over all investors as $\mathcal{G}$. This gives us the following:

$$\mathcal{G}^{\text{realized}} = \underbrace{\theta_G}_{\substack{\text{Each switcher} \\ \text{realizes this}}} \times \underbrace{[1 - F(r_H^*)]}_{\substack{\text{Fraction who} \\ \text{switch}}} + \underbrace{0}_{\substack{\text{Stayers realize} \\ \text{nothing}}} \times F(r_H^*) = \theta_G[1 - F(r_H^*)] \tag{30}$$

The total extra growth dividend to the government then simply $g\int_{r_H^*}^{\bar{r}}(r_H - r_L)f(r_H)dr_H$. Intuitively, we are taking the growth dividend g and growing it at an *additional* growth rate of

$r_H - r_L$, with r_H draw from the probability distribution function $f(r_H)$ with corresponding cumulative distribution function $F(r_H)$. Therefore, the total realization revenue is:

$$\mathcal{R}(\tau) = \tau \underbrace{\theta_G[1 - F(r_H^*)]}_{\substack{\text{Realization } revenue \\ G^{\text{realized}}}} + g \int_{r_H^*(\tau)}^{\bar{r}} (r_H - r_L) f(r_H) dr_H$$

This yields an interesting insight: From the integral, as τ increases, $r_H^*(\tau)$ increases. This means, that as τ increases, the proportion of investors who switch decreases. Thus, with a higher capital gains tax, there is less additional growth dividend to accrue from investors reallocating their capital into higher growth assets. The growth dividend therefore pushes the revenue-maximizing rate lower than would be the case without such a growth dividend.

Proposition 8 (Lock-in and the switching hurdle). Under a realization-based tax, an investor with embedded-gain fraction $\theta_G = \frac{G}{P}$ holding the low growth return asset (return is r_L) reallocates to a high-growth asset only if its return exceeds the hurdle $r_H^* = \frac{r_L}{1-\tau\theta_G}$. The hurdle is strictly increasing in both τ and θ_G: $\frac{dr_H^*}{d\tau} = \frac{r_L\theta_G}{(1-\tau\theta_G)^2} > 0$, $\frac{dr_H^*}{d\theta_G} = \frac{r_L\tau}{(1-\tau\theta_G)^2} > 0$. Long held, highly appreciated, positions are the most locked in. Uniqueness of the switching hurdle is in Appendix A.4.

We plot the lock-in effect in Figure 3. Here, we plot the hurdle rate, $\frac{r_H^*}{r_L}$ as a function of the tax rate τ for different levels of embedded gain θ_G. The key insights are: (a) the higher the tax rate the higher is the hurdle rate (i.e., required gain on the high growth asset) necessary to encourage switching. Intuitively, this is because the cost of selling an asset, then redeploying the capital, is higher for a higher tax rate. Further, (b), the higher the embedded gain, the higher the hurdle rate required for switching. Intuitively, this is because the investor bears a greater tax cost from the switch. The implication of the figure is that capital gains taxes can deter the efficient redeployment of capital.

[Insert Figure 3 about here]

Corollary 8.1 (Lock-in attenuates the growth dividend). Because $\frac{dr_H^*}{d\tau} > 0$, a higher rate lowers the share $1 - F(r_H^*)$ of investors who reallocate, reducing realized gains $\mathcal{G}$ and the realization revenue $\mathcal{R}(\tau)$. The realization margin therefore reinforces Proposition 7: Lock in effects push the revenue-maximizing rate further below 1/2.

6.2 What about investors who may never sell?

The next issue pertains to investors who may never sell. This has spiritual resonances of the setup in Chamley (1986). Intuitively, such a setup might apply to ultra high net worth individuals (UHNWIs) who have more assets than they need for consumption. It might also apply to families operating under a family office. With some modification, it can further apply to investors who borrow against assets rather than selling them. In this scenario, the investor has an additional choice: hold the asset forever.

To analyze this, let $\xi \in [0,1]$ be the probability that the gain entirely escapes capital taxation. This captures the possibility of capital taxation via restructuring needs, for example. Thus, $\xi = 1$ is certain escape from capital taxation, $\xi = 0$ means that the capital is taxed at the horizon regardless. Therefore, the effective holding tax is $\hat{\tau} = (1 - \xi)\tau$. There are therefore three terminal wealth effects. The hold then sell effect is $W_{\text{hold}} = (1 - \tau)P(1 + r_L) + \tau b$. If the investor switches investments, they can either sell at the end of the period, in which case their wealth is as before: $W_{\text{switch,sell}} = [P - \tau G][1 + (1-\tau) r_H]$. Alternatively, if they switch, then never sell, their wealth is $W_{\text{switch, don't sell}} = [P - \tau G][1 + (1-\hat{\tau}) r_H]$. Here, we have replaced one τ term with $\hat{\tau}$ to reflect the predicted tax on the asset held at maturity. If the investor never sells the asset, then we have $W_{\text{never sell}} = (1-\hat{\tau})P(1+r_L) + \hat{\tau} b$, which is simply the value of holding the low growth asset to maturity, incorporating the possibility that the wealth might be taxed regardless. If there is no tax at maturity, then the wealth becomes $P(1+r_L)$.

The foregoing setup gives us the new high growth rate $r_H^{\dagger}$ at which the investor might choose to switch. Recall that $G = \theta_G P$ and $b = P(1-\theta_G)$, where θ_G is the gain fraction $\frac{G}{P}$ with $G = P - b$. We now aim to determine the rate that renders it equivalent to switch as opposed to never sell (which offers a strictly better outcome than selling at the end of the second period). That is, we set $W_{\text{switch, don't sell}} = W_{\text{never sell}}$. We substitute G and b into the expressions for $W_{\text{switch, don't sell}}$ and $W_{\text{never sell}}$ and divide by P. This yields:

$$(1-\tau\theta_G)[1 + (1-\hat{\tau}) r_H] = (1-\hat{\tau})(1+r_L) + \hat{\tau}(1-\theta_G)$$

We can then collect terms and rearrange to obtain:

$$r_H^{\dagger} = \frac{r_L(1-\hat{\tau}) + \xi\tau\theta_G}{(1-\tau\theta_G)(1-\hat{\tau})} = \underbrace{\frac{r_L}{1-\tau\theta_G}}_{\text{baseline}} + \underbrace{\frac{\xi\tau\theta_G}{(1-\tau\theta_G)(1-\hat{\tau})}}_{\text{step-up premium}}$$

The interpretation is clear. We plot this in Figure 4. If the investor will always face a terminal tax rate, so $\xi = 0$, then the switching rate is the baseline rate: $r_H^{\dagger} = r_H^{*}$. By contrast, if the investor is certain to escape terminal tax, then $\xi = 1$ and $r_H^{\dagger} > r_H^{*}$. That is, the rate required for it to be worth switching assets is higher. This relationship increases with the level of embedded gain, θ_G, as illustrated in Figure 4.

[Insert Figure 4 about here]

Proposition 9 (Deferral and the step-up premium). Let $\xi \in [0,1]$ be the probability that a deferred gain escapes tax, so the effective holding tax is $\hat{\tau} = (1-\xi)\tau$. The switching hurdle for a deferral-capable investor is $r_H^{\dagger} = \frac{r_L(1-\hat{\tau})+\xi\tau\theta_G}{(1-\tau\theta_G)(1-\hat{\tau})} = r_H^{*} + \frac{\xi\tau\theta_G}{(1-\tau\theta_G)(1-\hat{\tau})}$. $r_H^{\dagger}$ is strictly increasing in ξ. Investors able to defer indefinitely (UHNW individuals, family offices, borrow-against-asset strategies) face a strictly higher switching threshold, intensifying lock-in.

7 How large are the magnitudes? An illustrative quantification

This section aims to partly quantify the effects of the foregoing discussion. While we do not claim to point-estimate the result from micro-data, we do aim to draw numerical insights from the model. We summarize the revenue-maximizing and welfare-maximizing capital gains tax rates, and their drivers, in Table 1.

[Insert Table 1 about here]

In our models, four main parameters are of interest. First, the sensitivity ν of investment effort to the net of tax rate. That is, how much does investment-effort (Ae^{*}) change with respect to the amount the investor gets to keep ($\beta = (1-\tau)$). This factors into the isoelastic cost function. It is reminiscent of the elasticity in the taxable income literature (Saez, 2001; Saez et al., 2012), and capital-related literature (Agersnap and Zidar, 2021). Second, the relative social value of a public dollar, λ, notably depends on the efficiency of government. This has resonances of the approach in Hendren and Sprung-Keyser (2020). This is important because maximizing *welfare* as opposed to maximizing government *revenue* is arguably a better goal, leaving aside the normative issue of whether the government *should* attempt to raise as much money as possible (as opposed to raising the minimum necessary). Third, the fiscal spillover parameter ϕ and the growth dividend g reflect the downstream impacts of capital. For our purposes, we focus on the tax rate's sensitivity to these and they have similar net effects on the optimal rate. Fourth, the substitution index, ρ_s, is anchored in the literature that people are mobile and that the wealthier, or more successful, is the person the more mobile they are. This is consistent with the literature that demonstrates mobility especially in 'superstars' (Akcigit et al., 2016; Kleven et al., 2013).

7.1 Revenue maximizing rates

We start by analyzing revenue maximizing rates. We emphasize that *revenue* maximization is myopic as it implies that the government spends money wisely, which is a strong assumption. An example from Australia, which has pushed to move CGT towards 47% of real gains, illustrates the point: in Victoria in Australia, anti-corruption investigations estimated that corruption associated with infrastructure builds cost taxpayers AUD 15 billion in a program worth around AUD 100 billion (Kolovos, 2026). Thus, in the presence of self-interested, or incompetent, administrators, welfare could arguably be maximized by keeping more money in private hands.

We plot the relationship between revenue-maximizing rates, elasticities (of investment effort to percentages retained), and fiscal spillovers in Figure 5. The figure presents some interesting findings. If there are no spillovers (i.e., $\phi = 0$) then the revenue-maximizing rate is strictly higher than if there are spillovers for a given level of keep rate elasticity. To put some numbers on this, we can get Australia's proposed 47% tax rate if there are no fiscal spillovers and keep rate elasticity $\nu = 1.13$. As will become apparent, such an assumption is unrealistic: the extant literature suggests that even absent fiscal spillovers, ν could be well into the 2s.

[Insert Figure 5 about here]

The question is then what is a realistic value of ν. This is an empirical question the depends on (a) how sensitive is investment effort to the amount of money the investor can retain, and (b) whether one incorporates fiscal spillovers into the analysis. There is significant variation in empirical 'elasticities', which themselves relate to the empirical Laffer curve peak. The literature tends to focus on the sensitivity of capital investment to tax rates, often omitting the downstream, spillover, or growth effects of capital investment. However, empirical Laffer curve peaks could implicitly reflect factors such as portfolio optimization strategies, such as strategic timing of realized gains and losses (Dai et al., 2008; Poterba and Weisbenner, 2001). Furthermore, our paper focuses on the tax on *real* (cf. nominal) gains whereas empirical Laffer peaks can relate to nominal gains. Nevertheless, they give some insight into potential parameter values.

We can nevertheless draw some insights in relation to potential elasticities. Dowd et al (2019) suggest a Laffer curve peak in the low-to-mid-20s, implying $\nu \approx 3$. Bakija and Gentry (2014) use state-level data and find a Laffer-peak in the high-20s, implying $\nu \approx 2.33$ at a peak of 30%. Agersnap and Zidar (2021) suggest the Laffer peak might be in the mid-30s, implying $\nu \approx 1.86$ at a peak of 35%. Given the significant variation in empirical estimates, and the fact that taxes are 'sticky' (i.e., policy mistakes are difficult to reverse), we contend that a lower Laffer peak is warranted. However, there is tension: realization-based Laffer estimates also embed the timing/lock-in margin we model separately in Section 6. Thus, they can over-state the structural effort-elasticity ν. Conservatively, we could consider the ν values we infer from empirical Laffer peaks as an upper bound. Furthermore, we note that such Laffer peaks typically ignore the fiscal spillover benefits of investing, so are most analogous to the case when we assume $\phi = 0$.

A related question is about the appropriate value of fiscal spillovers, ϕ and growth. Here, ϕ can be thought of as the additional fiscal benefits associated with that specific investment and g can be thought of as the additional growth that the investment sets in motion. The two factors are conceptually different, but have similar implications for optimal tax rates.

The revenue-maximizing tax rate is significantly higher if we ignore downstream effects to investment. However, the empirical literature shows significant economic benefits to investment (Dimitrova and Eswar, 2023). And, to be clear, the secondary investment market is necessary for an efficient and functional primary investment market, as is established in the capital raising literature. Nevertheless, appropriate values for ϕ might vary across investment-type. We can consider some anchors. The OECD average tax-to-GDP ratio is around 34% (OECD, 2025). This would understate the fiscal spillover from investing given that investing generates returns over future years (i.e., investing in a startup, which compounds, creates ongoing revenue).

We can also consider this from the perspective of the growth formula $g = \sum_{t \geq 1} \frac{(1+\gamma)^{t-1} \theta \beth}{(1+i)^t} = \frac{\theta \beth}{i - \gamma}$, where i is the cost of capital, γ the growth rate, θ the blended tax rate of the fiscal spillovers and $\beth$ the growth impact of a dollar of investing. One way to capture this is to analogize $\beth$ to the return on invested capital, which is estimated to be above 17%, albeit with significant time and industry variation (Damodaran, 2026a). But, that documented ROIC of 17% is after a US corporate tax rate of 21%, meaning $\beth$ in our formula would be approximately 21.5%. Therefore, if we assume θ is the corporate rate of 21%, $\beth = 21.5\%$, the GDP growth rate is 3% and the corporate cost of capital is 12.5%,[4] then we obtain $g \approx 0.475$. Arguably, such an estimate understates the growth-related benefits to investing as they focus on the ROIC of the investment to the firm per se, rather than considering broader societal benefits, including the ecosystem of firms, and innovation, that might accrue as a result of that investment.

If we incorporate potential estimates for the sensitivity of investment-effort to net-of-tax rates, and potential spillovers, we obtain interesting results. Having regard to Figure 5 and to the formula $\tau^R = \frac{1-\nu\phi}{1+\nu}$, we see that the revenue-maximizing rate is at most low and positive, tending towards zero. For example, taking low estimates of $\nu = 2$, and $\phi + g = 0.4$ we obtain $\tau^R = 0.07$. And, we note that both ν and $\phi + g$ could exceed those numbers, as indicated above.

The relationship between keep-rate elasticities, fiscal spillovers, and tax rates is interesting. Table 2 contains the revenue-maximizing tax rate (τ^R) for values of fiscal spillover (ϕ) and keep-rate elasticity (ν). To ground this in Australia's proposed 47% tax on real (cf nominal) realized capital gains, we can see that such a rate is supported only if we assume a low keep-rate elasticity or low fiscal spillovers. There is a region where total revenue is increased via an investment subsidy. Appendix A.5 notes the region boundaries. This arises where the fiscal spillover is large. Recall that we previously estimated an illustrative fiscal spillover of 0.475. At such a rate, a subsidy is unlikely to be optimal. However, the welfare-maximizing rate is likely low given the implied keep-rate elasticities (above) between 1.5 and 2.5.

[Insert Table 2 about here]

[4] We note there is significant variation in the implied equity risk premium; and thus, cost of capital across countries (Damodaran, 2026b; Fernandez et al., 2026).

7.2 Welfare maximizing rates

We next consider welfare maximizing rates. These differ from the revenue maximizing rate in that they also consider the value of keeping money in private hands. The welfare maximizing rate posits that it is not automatically optimal for the government to raise as much revenue as possible. This is especially pertinent if bureaucrats are corrupt, incompetent, or afflicted with agency conflicts. To use an extreme example: if a corrupt government were to siphon all tax revenue into waste or self-dealing, then the value of public funds would tend towards zero.

We start by exploring the relationship between welfare-maximizing rates and the social value of a public dollar. We plot this in Figure 6. For expositional clarity, we assume that there are no fiscal spillovers or growth effects in this graph (i.e., $\phi = g = 0$). The graph clearly shows that the welfare-maximizing rate is positive if and only if the government spends funds "appropriately". However, a corrupt or self-interested government might reduce the welfare-maximizing rate towards zero (or potentially into a subsidy, implying it is better to remove money from government).

[Insert Figure 6 about here]

We also consider the welfare-maximizing rate with respect to the elasticity of investment effort with respect to the keep rate (i.e., ν). We plot this in Figure 7. We plot this for various fiscal spillover levels (see Panels A-C). The key finding is that a tax rate of 47% (as Australia) proposes requires either a high value of public revenue (i.e., high λ) and/or a low elasticity of investment-to-keep-rate (i.e., ν). But, even if we assume $\lambda = 2$, and there are no fiscal spillovers, a 47% capital gains tax rate requires a keep-rate elasticity of approximately $\nu = 0.56$. But, as indicated, the keep-rate elasticity could well be closer to $\nu = 2$.

[Insert Figure 7 about here]

We present the welfare maximizing rates at various values of keep-rate elasticity (ν) and fiscal spillover (ϕ) in Table 3. In this table, we assume that the relative social value of a public dollar, λ, is 2, which reflects a strong preference for public funds. However, the valid rate depends on perceptions about public efficiency, which would vary across jurisdiction. Were the government to waste the money it derives from taxes, then λ should be low. There are several interesting results. First, the welfare-maximizing rate is lower than the revenue-maximizing rate unless we assume zero keep-rate elasticity. Furthermore, to ground this in real world numbers, at $\lambda = 2$, a 47% tax on real realized gains would assume a low keep rate elasticity, ν, which does not appear to be well supported in the literature. Second, if we use the fiscal spillover estimate (above) of approximately 0.475, we see that the welfare-maximizing tax rate tends towards zero.

[Insert Table 3 about here]

8 Conclusion

Government typically wants to encourage economic growth. Investment is a well-documented driver of, and prerequisite to, such growth. Thus, governments seek to encourage such investment. But, governments also frequently seek to tax the gains from investment. The question is then how best to tax such gains, if at all. This problem has resonances of an incentive contracting problem: the government looks to incentivize costly investment 'effort', which takes the form of capital, due diligence, and analysis. Thus, in contrast to extant literature, we recast the capital gains tax problem as one of moral hazard, akin to the problems discussed in Holmstrom (1979) and Holmstrom and Milgrom (1991, 1987).

Worked through, the framework delivers a consistent message. The naïve 'revenue maximizing' tax rate is myopic. Under stringent, and unrealistic, assumptions one can recover a tax revenue maximizing rate of one half of *real* (cf. *nominal*) gains. However, this does *not* maximize welfare. Furthermore, that revenue maximizing rate arises only if we ignore fiscal spillovers from investment, such as corporate tax, payroll tax, and the income tax attached to the jobs it creates. After incorporating such effects, the revenue-maximizing rate falls below one half, and beyond a threshold spillover the optimal policy inverts into a subsidy. Crediting the longer-run growth dividend pushes the rate lower again. Allowing investors to divert effort toward lower-taxed assets or jurisdictions drives the optimal rate on the taxed channel toward zero as the alternatives become close substitutes. Throughout, the welfare-maximizing rate sits below even the revenue-maximizing one, approaching one half only in the limiting case where the government values public funds infinitely more than private surplus.

Recognizing that the tax falls on realized rather than accruing gains sharpens the conclusion further. Because the liability is triggered only by sale, a higher rate raises the hurdle return at which an investor is willing to reallocate capital, locking holders into inferior assets and suppressing exactly the reallocation that generates the growth dividend. The effect is magnified for investors who can defer indefinitely. This includes ultra high net worth individuals, family offices, or individuals who can borrow against assets. In this case, the threshold at which an investor might be incentivized to switch from low, to high, growth assets increases. In other words, the higher the capital gains tax rate, the worse the lock-in effect. In turn, this undermines the aforementioned fiscal spillovers, and growth effects, the government might otherwise hope to accrue.

This paper does not analyze the normative arguments surrounding redistribution. Rather, we focus on the concern that CGT rates can reduce the tax base itself. We highlight that the high capital gains taxes do not automatically raise more revenue, and may reduce the tax base, which itself would undermine redistribution goals to the extent that redistribution is revenue-funded. Thus, even if the government's goal is to raise as much tax revenue as possible, simply raising capital rates is not always optimal, especially given the current capital gains tax rates in many advanced economies are above what our model suggests would be optimal.

The policy implications are direct. Governments must consider the growth impact of investment when considering capital taxation. Governments must also consider capital taxation in an international environment where international investment vehicles present as lower taxed alternatives. The naïve approach of pushing capital tax ever higher can undermine total tax revenue and reduce welfare. This has clear implications for countries that have moved to push capital taxes towards the top marginal tax rate: these run in the wrong direction if the objective is to maximize either welfare or revenue.

9 References

Agersnap, O., Zidar, O., 2021. The Tax Elasticity of Capital Gains and Revenue-Maximizing Rates. American Economic Review: Insights 3.

Aguiar, M.A., Moll, B., Scheuer, F., 2024. Putting the “Finance” into “Public Finance”: A Theory of Capital Gains Taxation (NBER Working Paper No. 32951). https://doi.org/10.3386/w32951

Agyemang, E., 2025. Capital gains tax receipts fall after big cuts in allowances - Revenue falls 18 per cent in 2023-24 to £12.1bn as the UK’s efforts to increase revenues from the levy backfire. Financial Times.

Akcigit, U., Baslandze, S., Stantcheva, S., 2016. Taxation and the International Mobility of Inventors. American Economic Review 106, 2930–2981. https://doi.org/10.1257/aer.20150237

Arefeva, A., Davis, M.A., Ghent, A.C., Park, M., 2025. The Effect of Capital Gains Taxes on Business Creation and Employment: The Case of Opportunity Zones. Management Science 71, 4566–4581. https://doi.org/10.1287/mnsc.2022.03223

Ayers, B.C., Lefanowicz, C.E., Robinson, J.R., 2003. Shareholder Taxes in Acquisition Premiums: The Effect of Capital Gains Taxation. Journal of Finance 58, 2783–2801. https://doi.org/10.1046/j.1540-6261.2003.00622.x

Bakija, J.M., Gentry, W.M., 2014. Capital gains taxes and realizations: Evidence from a long panel of state-level data (Working paper). Williams College.

Barnett, J., 2025. Capital gains tax receipts fall 10% as wealthy exit UK. The Times.

Benedetti-Fasil, C., Sedlacek, P., Sterk, V., 2022. Startups and employment following the COVID-19 pandemic: a calculator. Economic Policy 37, 507–533. https://doi.org/10.1093/epolic/eiac028

Bernile, G., Cumming, D., Lyandres, E., 2007. The size of venture capital and private equity fund portfolios. Journal of Corporate Finance 13, 564–590.

Borensztein, E., De Gregorio, J., Lee, J.-W., 1998. How does foreign direct investment affect economic growth? Journal of International Economics 45, 115–135. https://doi.org/10.1016/S0022-1996(97)00033-0

Campello, M., Junqueira, G., 2025. Tax Incentives and Venture Capital Risk-Taking (Working Paper). University of Florida.

Chamley, C., 1986. Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives. Econometrica 54, 607–622. https://doi.org/10.2307/1911310

Chen, J., Hshieh, S., 2024. The labor effects of R&D tax incentives: evidence from VC-backed startups. Review of Finance 28, 1451–1482. https://doi.org/10.1093/rof/rfae026

Chetty, R., 2009. Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods. Annual Review of Economics 1, 451–488. https://doi.org/10.1146/annurev.economics.050708.142910

Chow, T., Huang, S., Klassen, K.J., Ng, J., 2022. The Influence of Corporate Income Taxes on Investment Location: Evidence from Corporate Headquarters Relocations. Management Science 68, 1404–1425. https://doi.org/10.1287/mnsc.2020.3906

Cullen, J., Gordon, R., 2007. Taxes and entrepreneurial risk-taking: Theory and evidence for the U.S. Journal of Public Economics 91, 1495–1505.

Cumming, D., Dai, N., 2011. Fund size, limited attention and valuation of venture capital backed firms. Journal of Empirical Finance 18, 2–15. https://doi.org/10.1016/j.jempfin.2010.09.002

Dai, Z., Maydew, E., Shackelford, D.A., Zhang, H.H., 2008. Capital Gains Taxes and Asset Prices: Capitalization or Lock-in? Journal of Finance 63, 709–742. https://doi.org/10.1111/j.1540-6261.2008.01329.x

Damodaran, A., 2026a. Margin/ ROIC by Sector (US) [WWW Document]. URL https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/mgnroc.html (accessed 6.21.26).

Damodaran, A., 2026b. Equity Risk Premiums (ERP): Determinants, Estimates and Implications - The 2026 Edition (Working Paper). NYU Stern School of Business. https://doi.org/10.2139/ssrn.6361419

de Montaigne, M., 1580. That Men by various Ways arrive at the same end, in: Screech, M.A. (Trans.), The Complete Essays. Penguin.

Dimitrova, L., Eswar, S.K., 2023. Capital Gains Tax, Venture Capital, and Innovation in Start-Ups. Review of Finance 27, 1471–1519. https://doi.org/10.1093/rof/rfac057

Domar, E.D., Musgrave, R.A., 1944. Proportional Income Taxation and Risk-Taking. Quarterly Journal of Economics 58, 388–422. https://doi.org/10.2307/1882847

Dowd, T., McClelland, R., Mortenson, J.A., 2019. Investor Responsiveness to Capital Gains Taxes During the Great Recession. Urban Institute.

Edwards, A., Todtenhaupt, M., 2020. Capital gains taxation and funding for start-ups. Journal of Financial Economics 138, 549–571. https://doi.org/10.1016/j.jfineco.2020.06.009

Ehling, P., Tompaidis, S., Yang, C., 2026. Tax revenue from realized capital gains. Review of Finance 30, 863–886. https://doi.org/10.1093/rof/rfaf078

Feldstein, M., 1999. Tax Avoidance and the Deadweight Loss of the Income Tax. Review of Economics and Statistics 81, 674–680.

Feldstein, M., 1995. The Effect of Marginal Tax Rates on Taxable Income: A Panel Study of the 1986 Tax Reform Act. Journal of Political Economy 103, 551–572.

Fernandez, P., Habibian, A., Fdez Acin, A., 2026. Survey: Market Risk Premium and Risk-Free Rate used for 97 countries in 2026 (Working paper). https://doi.org/10.2139/ssrn.6570342

Fischer, M., Gallmeyer, M.F., 2016. Heuristic portfolio trading rules with capital gain taxes. Journal of Financial Economics 119, 611–625. https://doi.org/10.1016/j.jfineco.2016.01.024

Gentry, W.M., Hubbard, R.G., 2000. Tax Policy and Entrepreneurial Entry. American Economic Review 90, 283–287. https://doi.org/10.1257/aer.90.2.283

Golosov, M., Kocherlakota, N., Tsyvinski, A., 2003. Optimal Indirect and Capital Taxation. Review of Economic Studies 70, 569–587.

Gordon, R., Sarada, 2018. How should taxes be designed to encourage entrepreneurship? Journal of Public Economics 166, 1–11. https://doi.org/10.1016/j.jpubeco.2018.08.003

Guenther, D.A., Willenborg, M., 1999. Capital gains tax rates and the cost of capital for small business: evidence from the IPO market. Journal of Financial Economics 53, 385–408. https://doi.org/10.1016/S0304-405X(99)00026-4

Hendren, N., Sprung-Keyser, B., 2020. A Unified Welfare Analysis of Government Policies. Quarterly Journal of Economics 135, 1209–1318. https://doi.org/10.1093/qje/qjaa006

Holden, R., 2026. There is no economic case for taxing work and investment the same. Australian Financial Review.

Holmstrom, B., 1979. Moral Hazard and Observability. Bell Journal of Economics 10, 74–91. https://doi.org/10.2307/3003320

Holmstrom, B., Milgrom, P., 1991. Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design. Journal of Law, Economics & Organization 7, 24–52.

Holmstrom, B., Milgrom, P., 1987. Aggregation and Linearity in the Provision of Intertemporal Incentives. Econometrica 55, 303–328. https://doi.org/10.2307/1913238

Huizinga, H., Voget, J., Wagner, W., 2018. Capital gains taxation and the cost of capital: Evidence from unanticipated cross-border transfers of tax base. Journal of Financial Economics 129, 306–328. https://doi.org/10.1016/j.jfineco.2018.04.014

Humphery-Jenner, M., 2013. Diversification in Private Equity Funds: On Knowledge-Sharing, Risk-Aversion and Limited-Attention. Journal of Financial and Quantitative Analysis 48, 1545–1572.

Humphery-Jenner, M., 2012. Private Equity Fund Size, Investment Size, and Value Creation. Review of Finance 16, 799–835. https://doi.org/10.1093/rof/RFR011

Jin, L., 2006. Capital Gains Tax Overhang and Price Pressure. Journal of Finance 61, 1399–1431. https://doi.org/10.1111/j.1540-6261.2006.00876.x

Judd, K.L., 1985. Redistributive taxation in a simple perfect foresight model. Journal of Public Economics 28, 59–83. https://doi.org/10.1016/0047-2727(85)90020-9

Keuschnigg, C., 2004. Taxation of a venture capitalist with a portfolio of firms. Oxford Economic Papers 56, 285–306. https://doi.org/10.1093/oep/gpf041

Keuschnigg, C., Nielsen, S.B., 2004. Start-ups, venture capitalists, and the capital gains tax. Journal of Public Economics 88, 1011–1042.

Keuschnigg, C., Nielsen, S.B., 2003. Tax policy, venture capital, and entrepreneurship. Journal of Public Economics 87, 175–203. https://doi.org/10.1016/S0047-2727(01)00170-0

Kleven, H.J., Landais, C., Saez, E., 2013. Taxation and International Migration of Superstars: Evidence from the European Football Market. American Economic Review 103, 1892–1924. https://doi.org/10.1257/aer.103.5.1892

Knill, A., 2009. Should venture capitalists put all their eggs in one basket? Diversification versus pure-play strategies in venture capital. Financial Management 38, 441–486.

Kocherlakota, N.R., 2005. Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation. Econometrica 73, 1587–1621. https://doi.org/10.1111/j.1468-0262.2005.00630.x

Kolovos, B., 2026. The CFMEU has cost Victorian taxpayers $15bn, a corruption fighter claims. How did he reach that figure and what happens next? The Guardian.

Lindsey, L., Stein, L.C.D., 2026. Angels, entrepreneurship, and employment dynamics: Evidence from investor accreditation rules. Journal of Financial Economics 183, 104323. https://doi.org/10.1016/j.jfineco.2026.104323

Lopez-de-Silanes, F., Phalippou, L., Gottschalg, O., 2015. Giants at the Gate: Investment Returns and Diseconomies of Scale in Private Equity. Journal of Financial and Quantitative Analysis 50, 377–411.

Mauer, D.C., Senbet, L.W., 1992. The Effect of the Secondary Market on the Pricing of Initial Public Offerings: Theory and Evidence. Journal of Financial and Quantitative Analysis 27, 55–79. https://doi.org/10.2307/2331298

Mirrlees, J.A., 1971. An Exploration in the Theory of Optimum Income Taxation. Review of Economic Studies 38, 175–208. https://doi.org/10.2307/2296779

OECD, 2025. Revenue Statistics 2025.

Poterba, J.M., Weisbenner, S.J., 2001. Capital Gains Tax Rules, Tax-loss Trading, and Turn-of-the-year Returns. Journal of Finance 56, 353–368. https://doi.org/10.1111/0022-1082.00328

Saez, E., 2001. Using Elasticities to Derive Optimal Income Tax Rates. Review of Economic Studies 68, 205–229. https://doi.org/10.1111/1467-937X.00166

Saez, E., Slemrod, J., Giertz, S.H., 2012. The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review. Journal of Economic Literature 50, 3–50.

Scheuer, F., 2014. Entrepreneurial Taxation with Endogenous Entry. American Economic Journal: Economic Policy 6, 126–163. https://doi.org/10.1257/pol.6.2.126

Sedlacek, P., Sterk, V., 2019. Reviving American entrepreneurship? Tax reform and business dynamism. Journal of Monetary Economics 105, 94–108.

Stewart, M., 2026. Submission: Inquiry into Treasury Laws Amendment (Tax Reform No. 1) Bill 2026 and a related bill (Senate Inquiry Submission).

Straub, L., Werning, I., 2020. Positive Long-Run Capital Taxation: Chamley-Judd Revisited. American Economic Review 110, 86–119. https://doi.org/10.1257/aer.20150210

Strulik, H., Trimborn, T., 2012. Laffer strikes again: Dynamic scoring of capital taxes. European Economic Review 56, 1180–1199. https://doi.org/10.1016/j.euroecorev.2012.05.002

Xu, J., 2023. Is There a Trade-off Between Protecting Investors and Promoting Entrepreneurial Activity? Evidence from Angel Financing. Journal of Financial and Quantitative Analysis 58, 3305–3341. https://doi.org/10.1017/S0022109022001314

10Tables

Table 1: Summary of key results in the paper

This table summarizes the key results in the paper. All rates apply to real gains and set the risk term $s = 0$ for exposition. λ is the relative social value of a public dollar; $\rho_s = \kappa/k \in [0,1)$ the substitution index; $\delta \in (0,1]$ the value of present consumption relative to investing (Prop. 6 requires $\rho_s < \sqrt{\delta}$); $\phi = \psi/\rho$ the capitalized fiscal spillover with ψ being the per period ancillary yield and ρ is the discount rate; $g = \sum_{t\geq 1} \frac{(1+\gamma)^{t-1}\theta ב}{(1+i)^t} = \frac{\theta ב}{i-\gamma}$ is the growth dividend, with i being the opportunity cost of capital, γ the growth rate, and ב the output and θ the tax rate applicable to it. The baseline and spillover rows generalize to any base elasticity ν by replacing ½ with $1/(1+\nu)$ and $2\lambda - 1$ with $\lambda(1+\nu) - 1$. The final two rows have no single closed-form rate because the realization basis acts on the reallocation margin through the hurdle return.

Model feature added (Prop.)	**Revenue-maximizing rate, $\boldsymbol{\tau^R}$**	**Welfare-maximizing rate, $\boldsymbol{\tau^*}$**	**Effect on the rate / key comparative static**
Baseline: single asset, symmetric losses (Prop. 1)	1/2	$\frac{\lambda - 1}{2\lambda - 1}$ Where λ is the relative social value of a public dollar in $W(\tau) = CE^*(\tau) + \lambda R(\tau)$ where $CE^*(\tau)$ is the private certainty equivalent and $R(\tau)$ is government revenue.	Anchor and ceiling: $\tau^* \leq \tau^R$ for all λ. $\tau^* \to 0$ as $\lambda \to 1$ and $\tau^* \to \frac{1}{2}$ only as $\lambda \to \infty$
General base elasticity (Prop. 2)	$\frac{1}{1+\nu}$ Where ν is the elasticity of investment to keep rate the cost function $C(e) = \frac{ke^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}}$	$\frac{\lambda - 1}{\lambda(1+\nu) - 1}$	Both fall as investment becomes more tax-responsive (i.e., $\nu \uparrow$); recover baseline at $\nu = 1$.
Imperfect loss offset (Prop. 3)	$\min\left\{\frac{1}{2}, \tau_{\text{participation}}(\ell)\right\}$	via participation	Worse offset ($\ell \downarrow$) lowers the exit threshold, capping the feasible rate below ½.
Fiscal spillover (Prop. 4)	$\frac{1-\phi}{2}$ Where ϕ is the ancillary revenue parameter with ancillary revenue $\phi E[Y]$	$\frac{\lambda - 1 - \lambda\phi}{2\lambda - 1}$	Each unit of spillover lowers both rates; policy flips to a subsidy once $\phi > \frac{\lambda-1}{\lambda}$
Fiscal spillover with generalized elasticity of investment to	$\frac{1-\nu\phi}{1+\nu}$	$\frac{(\lambda-1) - \lambda\nu\phi}{\lambda(1+\nu) - 1}$	

Model feature added (Prop.)	Revenue-maximizing rate, τ^R	Welfare-maximizing rate, τ^*	Effect on the rate / key comparative static
keep rate (Cor. 4.2)			
Lower-/untaxed alternative (Prop. 5)	$\frac{1-\rho_s}{2}$ Where ρ_s captures the substitutability between assets	$\frac{(1-\rho_s)(\lambda-1)}{2\lambda-1}$	Falls in substitutability ρ_s; → 0 as ρ_s → 1 (frictionless mobility). Uniform taxation across alternatives (Cor. 5.2) restores ½
Consumption margin (Prop. 6)	$\frac{1}{2}\left[1-\phi-\rho_s\left((1-\tau_c)-\frac{\tau_c}{\delta}\right)\right]$ Where τ_c is the consumption tax rate, δ reflects the relative value of consumption to investing	$\frac{(\lambda-1)[1-\rho_s(1-\tau_c)]+\frac{\rho_s\lambda\tau_c}{\delta}-\lambda\phi}{2\lambda-1}$	Higher consumption tax τ_c *raises* the rate (recapture); greater substitutability ρ_s or pull-to-consume δ lowers it
Growth dividend (Prop. 7)	$\frac{1}{2}\left[1-g-\rho_s((1-\tau_c)-\frac{\tau_c}{\delta})\right]$ Where g denotes the downstream growth impacts of the investment	$\frac{(\lambda-1)[1-\rho_s(1-\tau_c)]+\frac{\rho_s\lambda\tau_c}{\delta}}{2\lambda-1}-\frac{\lambda g}{2\lambda-1}$	g (present value of future tax stream) enters exactly as ϕ: crediting growth pushes the rate down again
Realization lock-in (Prop. 8 / Cor. 8.1)	Operates through the switching hurdle $r_H^*=\frac{r_L}{1-\tau\theta_G}$, where r_H is the return on the high growth asset, r_L is the return on the low growth asset and $\theta_G=\frac{G}{P}$, where P is the asset's value and G the embedded gain.	N/A	A higher rate raises the hurdle, freezes capital in inferior assets, shrinks realized gains, pushing τ^R lower
Deferral / step-up, when investors may never sell (Prop. 9)	Lock-in hurdle rate becomes $r_H^\dagger=\frac{r_L(1-\hat{\tau})+\xi\tau\theta_G}{(1-\tau\theta_G)(1-\hat{\tau})}$ Where ξ is the probability that the gain is never taxed and $\hat{\tau}=(1-\xi)\tau$ is the effective tax rate	N/A	Buy-borrow-die widens the hurdle further; lock-in is worst for UHNW investors, family offices, asset-backed borrowers

Table 2: Revenue-maximizing tax rates by keep-rate elasticity and fiscal spillover

This table contains figures for the revenue-maximizing tax rate $\tau^R = \frac{1-\nu\phi}{1+\nu}$ for values of fiscal spillover ϕ and keep-rate elasticity ν. The column heading contains the fiscal spillover value and the left-hand row heading contains the keep-rate elasticity.

		Fiscal Spillover				
		0	0.25	0.5	0.75	1
Keep-rate elasticity	0	100%	100%	100%	100%	100%
	0.5	67%	58%	50%	42%	33%
	1	50%	38%	25%	13%	0%
	1.5	40%	25%	10%	-5%	-20%
	2	33%	17%	0%	-17%	-33%
	2.5	29%	11%	-7%	-25%	-43%
	3	25%	6%	-13%	-31%	-50%

Table 3: Welfare-maximizing tax rates by keep-rate elasticity and fiscal spillover

This table contains figures for the welfare-maximizing tax rate $\tau^* = \frac{(\lambda-1)-\lambda\nu\phi}{\lambda(1+\nu)-1}$ for values of fiscal spillover ϕ and keep-rate elasticity ν. This table assumes that the relative social value of a public dollar, λ, is 2. The column heading contains the fiscal spillover value and the left hand row heading contains the keep-rate elasticity.

		Fiscal Spillover				
		0	0.25	0.5	0.75	1
Keep-rate elasticity	0	100%	100%	100%	100%	100%
	0.5	50%	38%	25%	13%	0%
	1	33%	17%	0%	-17%	-33%
	1.5	25%	6%	-13%	-31%	-50%
	2	20%	0%	-20%	-40%	-60%
	2.5	17%	-4%	-25%	-46%	-67%
	3	14%	-7%	-29%	-50%	-71%

11 Figures

Figure 1: Welfare maximizing rate and the social value of a public dollar

This figure plots the relationship between the welfare maximizing tax rate, τ^*, as a function of the perceived relative social value of a public dollar, λ. Here, per Proposition 1, the welfare maximizing rate $\tau^* = \frac{\lambda-1}{2\lambda-1}$.

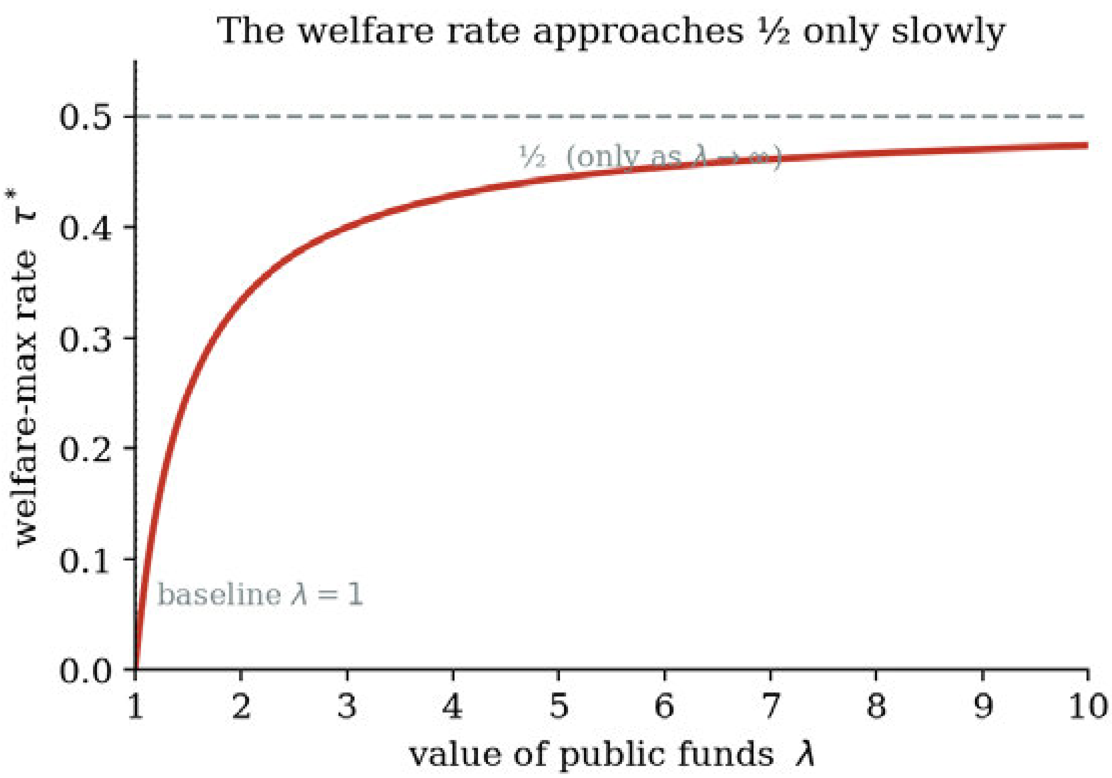

Figure 2: Loss offsets

This figure plots the tax rate below which investors participate (i.e., $\tau_{\text{participation}}(\ell)$) as a function of the loss-offset ratio, ℓ, for two values of the outside option $\bar{u}$ at $m = 1, s = 0.25$.

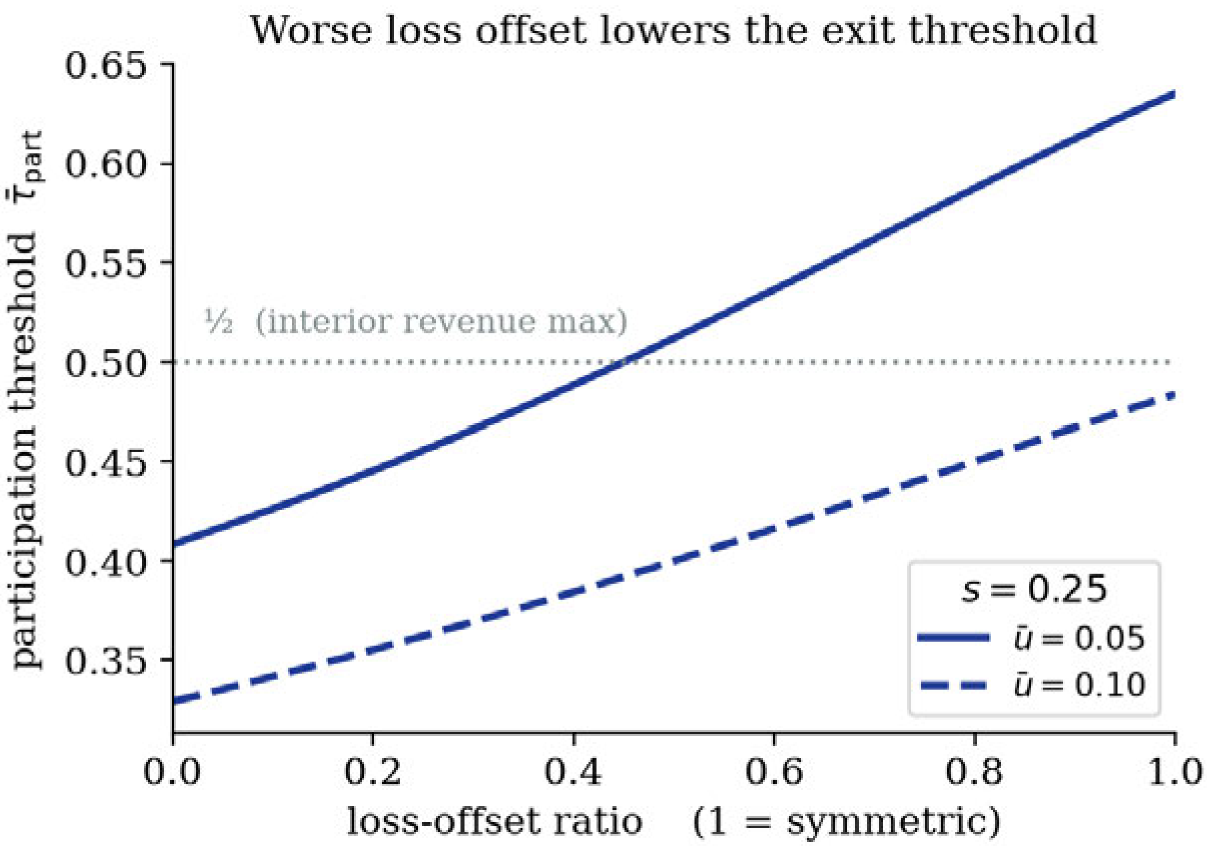

Figure 3: Hurdle rates to facilitate switching

This figure plots the hurdle rate, $\frac{r_H^*}{r_L}$, required to switch from the low growth asset to the high growth asset as a function of the capital gains tax rate τ. We plot it for various levels of embedded gain $\theta_G = \frac{G}{P}$, where G is the embedded capital gain, and P is the value of the asset. We obtain this from the relationship $r_H^* = \frac{r_L}{1-\tau\theta_G}$ as noted in Proposition 8.

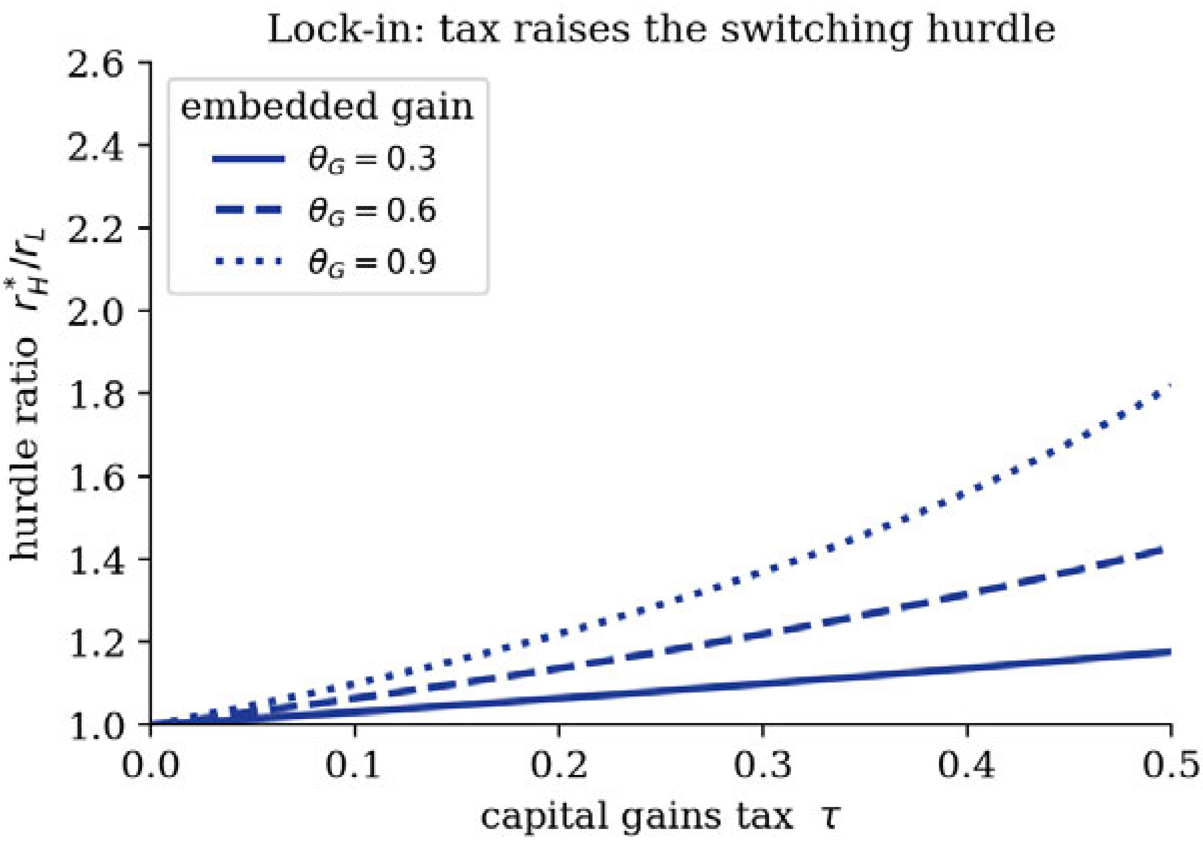

Figure 4: Required hurdle rates and the possibility of infinite investment horizons

This figure plots the required hurdle rate to switch investments, $r_H^\dagger/r_L$, as a function of the probability of being able to hold forever, not paying capital gains tax. We assume a capital gains tax in this graph of 35% and plot the relationship for various levels of embedded gain θ_G.

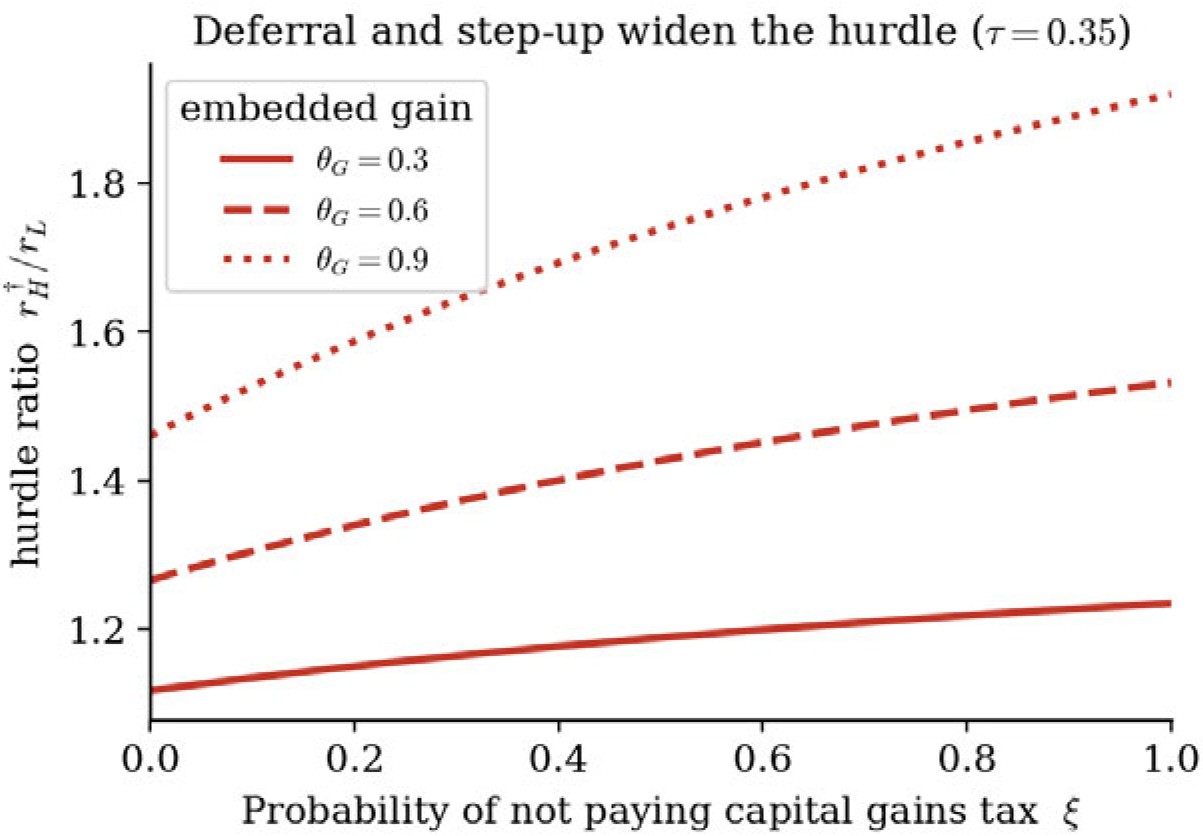

Figure 5: Revenue-maximizing rates, elasticity, and fiscal spillovers

The revenue-maximizing rate and proposed policy rates. The figure plots the revenue-maximizing rate $\tau\hat{}R\_\phi = (1 - \nu\phi)/(1 + \nu)$ (Corollary 4.2) against the base elasticity ν, for various fiscal-spillover levels. Horizontal lines mark Australia's proposed 47% rate and the UK's 24% rate. The marked point shows that, absent spillover, 47% is revenue-maximizing only for $\nu = 1.13$. But, the maximum elasticity consistent with a 47% revenue-maximizing rate falls when spillovers are introduced.

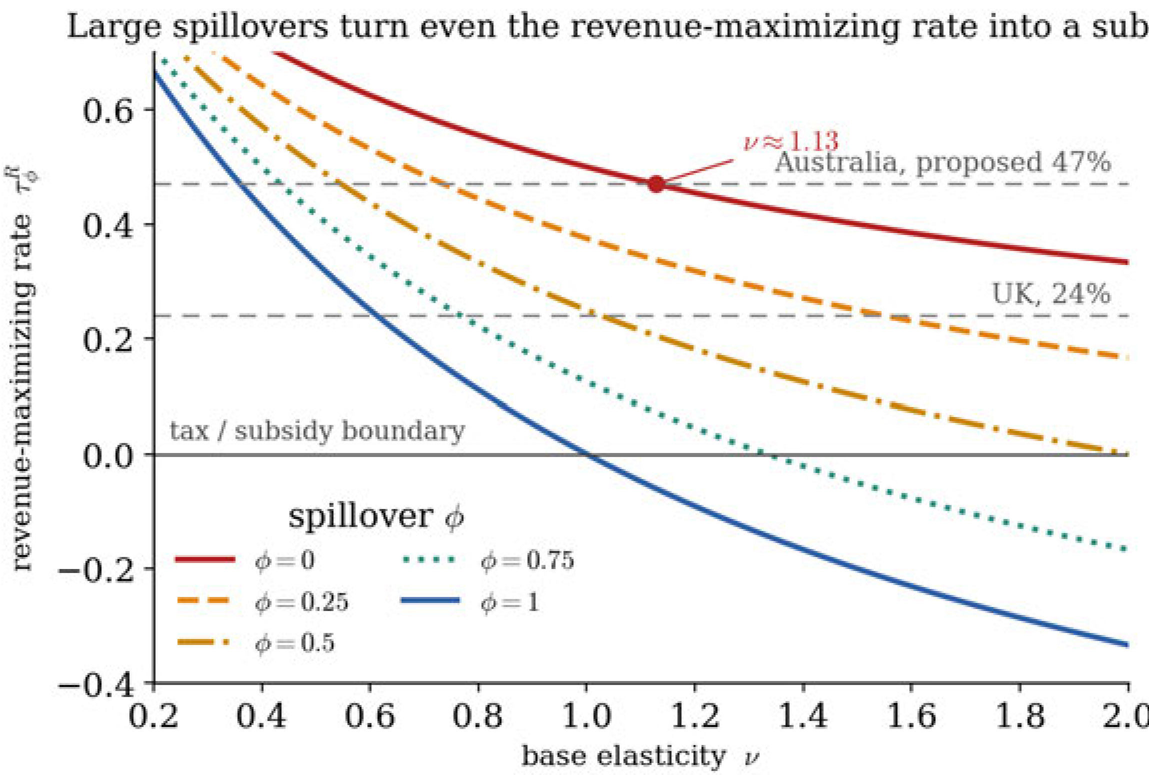

Figure 6: Welfare maximizing rates

This figure plots the welfare maximizing tax rate τ^* as a function of the 'elasticity' or 'sensitivity' parameter ν (the sensitivity of investment-effort to the net of tax rate $(1-\tau)$, and the relative social value of a public dollar, λ, in $W(\tau) = CE^*(\tau) + \lambda R(\tau)$. For expositional clarity, this graph assumes that there is no fiscal spillover.

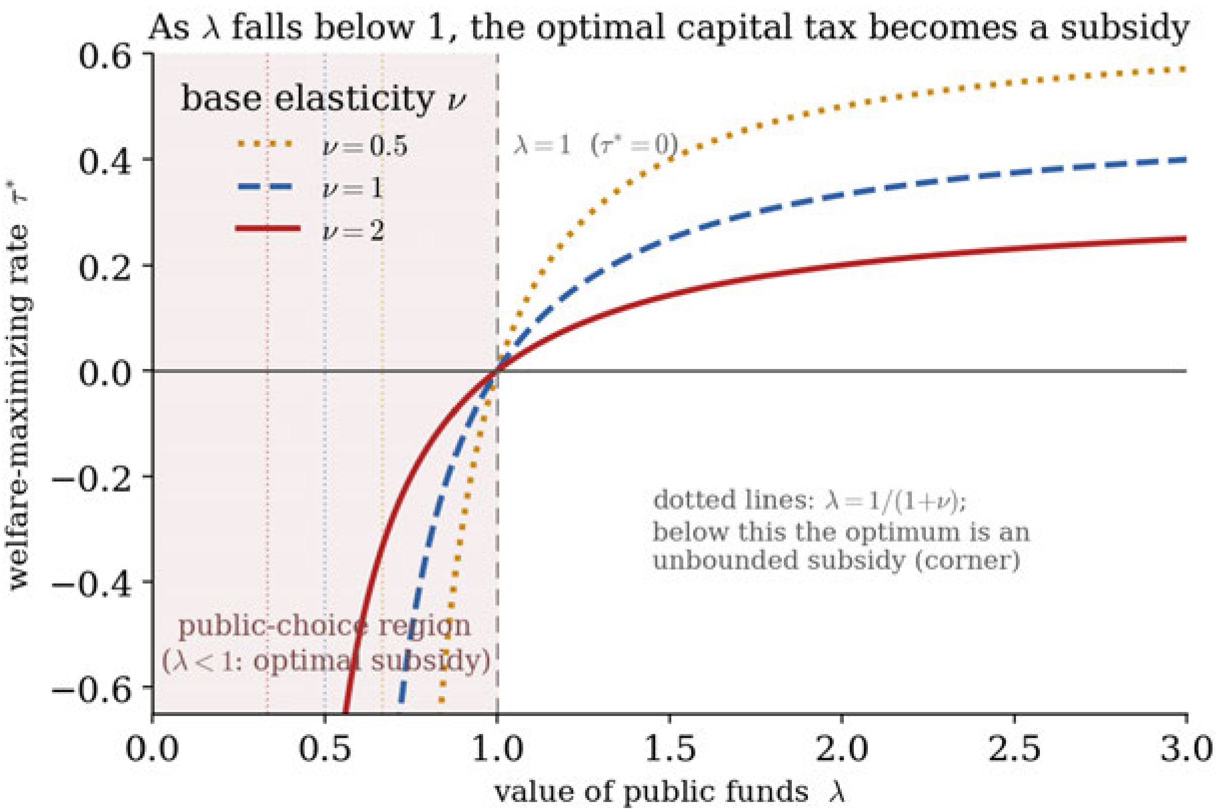

Figure 7: Welfare maximizing rates, social value of a public dollar, and elasticity

This figure plots the welfare maximizing tax rate, τ^* for various levels of keep rate elasticity, ν, at various levels of fiscal spillover ϕ. This arises from the formula $\tau^* = \frac{(\lambda-1)-\lambda\nu\phi}{\lambda(1+\nu)-1}$.

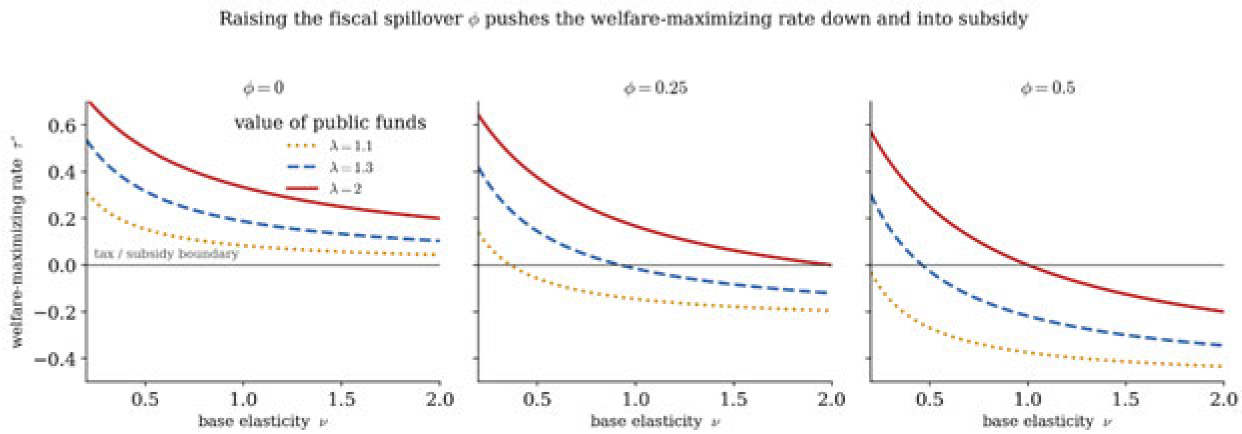

Appendix A: Second order conditions, concavity, regime boundaries

This appendix collects the second-order conditions (SOCs) and concavity arguments that establish the optima reported in the main text as genuine maxima, the positive-definiteness conditions underlying the multitask problems, and the boundaries separating interior-tax, zero, and subsidy regimes. Throughout we use the main-text notation: $\beta = 1 - \tau$, $m = A^2/k$, $s = a\sigma^2/m$, and $W(\tau) = CE*(\tau) + \lambda R(\tau)$. Unless noted, the risk term is retained (results with $s = 0$ follow by setting $s = 0$).

A.1 The investor's effort problem (Lemmas 1–2)

The investor chooses effort to maximize the certainty equivalent. We verify concavity for both cost specifications, which establishes that the effort rule e*(τ) in Lemma 2 is the unique maximizer.

Quadratic cost: With $C(e) = \left(\frac{k}{2}\right)e^2$, the certainty equivalent $CE(e) = \beta Ae - \left(\frac{k}{2}\right)e^2 - \left(\frac{a}{2}\right)\beta^2\sigma^2$ has

$$CE'(e) = \beta A - ke, \quad CE''(e) = -k < 0.$$

The objective is strictly concave in e, so the first-order condition $\beta A - ke = 0$ delivers the unique global maximizer $e^*(\tau) = \beta A/k$. This is the SOC underlying Lemma 2 in the quadratic case.

Isoelastic cost: With $C(e) = k \cdot e\text{^}(1 + 1/\nu)/(1 + 1/\nu), \nu > 0$, the marginal cost and its derivative are

$$C'(e) = ke^{\frac{1}{\nu}}, \quad C''(e) = \left(\frac{k}{\nu}\right)e^{\frac{1}{\nu}-1} > 0 \ \ for\, e > 0$$

The agent's objective $\beta Ae - C(e)$ therefore has second derivative $-C''(e) < 0$ on $e > 0$: it is strictly concave, and the effort condition $\beta A = C'(e)$ yields the unique maximizer $e^*(\beta) = \left(\frac{\beta A}{k}\right)^{\nu}$ reported in equation (8). Hence Lemma 2's effort rule and the base elasticity d ln E[Y]/d ln β = v are properly grounded for all v > 0.

Lemma 1 (certainty equivalent of a CARA-normal payoff) is a standard property of the normal moment generating function. The proof is standard, but repeated for expositional clarity. For $x \sim N(\mu, v)$, the moment generating function gives $E[\exp(-ax)] = \exp\left(-a\mu + \frac{a^2}{2}v\right)$. Thus, $E[U(x)] = -\exp\left(-a\mu + \frac{a^2}{2}v\right) = -\exp\left(-a\left[\mu - \frac{a}{2}v\right]\right)$. However, $-\exp(-a)$ is strictly increasing. Therefore, maximizing $E[U]$ is the same as maximizing $\mu - \frac{a}{2}v$.

A.2 The government's single-channel problem (Propositions 1, 2, 4, 5, 7)

Revenue, baseline (Prop. 1): With quadratic cost, $R(\tau) = \tau(1-\tau)m$, so

$$R'(\tau) = (1 - 2\tau)m, \quad R''(\tau) = -2m < 0.$$

R is strictly concave; $\tau^R = \frac{1}{2}$ is the unique maximizer.

Revenue, general elasticity (Prop. 2). With isoelastic cost, $R(\tau) = m_\nu \tau(1-\tau)^\nu$, where $m_\nu = A\left(\frac{A}{k}\right)^\nu > 0$. Then

$$R'(\tau) = m_\nu(1-\tau)^{\nu-1} \cdot [(1 - \tau) - \nu\tau]$$

On [0,1) the bracket $(1 - \tau) - \nu\tau = 1 - (1+\nu)\tau$ is strictly decreasing and crosses zero once, at $\tau^R = \frac{1}{1+\nu}$; R' changes sign from positive to negative there. Thus R is single-peaked (strictly quasi-concave) on $[0,1)$ with unique maximizer $\tau^R = \frac{1}{1+\nu}$. (Note R is not globally concave for all ν, so single-peakedness (not $R'' < 0$ everywhere)

Welfare (Prop. 1). Differentiating $W(\tau) = \frac{(1-\tau)^2 m(1-s)}{2} + \lambda\tau(1-\tau)m$ and dividing by m,

$$\frac{W'(\tau)}{m} = -(1-\tau)(1-s) + \lambda(1 - 2\tau), \quad \frac{W''(\tau)}{m} = (1 - s) - 2\lambda$$

The interior stationary point is a maximum if and only if $W'' < 0$, i.e.

$$\lambda > \frac{1-s}{2}$$

This requires $\lambda > \frac{1-s}{2}$. Since $\frac{1-s}{2} < \frac{1}{2}$, this permits $\lambda < 1$. A value of $\lambda < 1$ would suggest government waste or self-dealing, as alluded to in the public-choice literature. In which case, a tax subsidy may be optimal, as explored in Section 7 and Figure 6. This fails only if $\lambda \le \frac{1-s}{2}$, where the welfare problem has a corner (subsidy) solution rather than an interior maximum.

Welfare (Prop. 2, general elasticity). The same computation with isoelastic cost replaces the SOC threshold by

$$W'' < 0 \;\Leftrightarrow\; \lambda(1+\nu) - 1 > 0,\; i.e.\; \lambda > \frac{1}{1+\nu}$$

which again holds for all $\lambda > \frac{1}{1+\nu}$ and permits $\lambda < 1$. The interior welfare rate $\tau^* = \frac{\lambda-1}{\lambda(1+\nu)-1}$ is therefore a maximum.

Spillover and growth-dividend extensions (Props. 4, 5, 7). The fiscal-spillover term ϕ and the growth dividend g enter government revenue additively and linearly in τ: $G(\tau) = (\tau + \phi)(1-\tau)m$, where $G(\tau)$ reflects the total of the direct tax revenue and the spillover. This is similar when incorporating growth effects, g. These effectively shift W' by a constant and leave W'' unchanged. Consequently the second-order conditions established above for Propositions 1–2 carry over unchanged to Propositions 4, 5, and 7; no separate SOC is required for the spillover- or growth-augmented rates. The same applies to the lower-/untaxed-alternative rates of Proposition 5, which scale the objective by the fixed positive factor (see A.3).

A.3 The multitask and consumption-margin problems (Section 4, Proposition 6)

In the multitask settings the agent's SOC is positive-definiteness of the effort cost Hessian, which guarantees the joint effort problem is strictly concave and the stationary point $\boldsymbol{e}^* = \boldsymbol{K}^{-1}\boldsymbol{b}$ is the unique maximizer.

Two-channel allocation (Section 4). With joint cost $C(e) = \frac{1}{2}\boldsymbol{e}^{\mathrm{T}}\boldsymbol{K}\boldsymbol{e}$. The agent's objective $\Sigma\, \beta_{\mathrm{i}} A_{\mathrm{i}} e_{\mathrm{i}} - C(e)$ has Hessian K. Strict concavity requires $K \succ 0$, i.e. $k_1 > 0$ and $det\, K =$

$k_1k_2 - \kappa^2 > 0$. In the symmetric case $k_1 = k_2 = k$ this is $k > 0$ and $k^2 - \kappa^2 > 0$, equivalently the substitution index $\rho_s = \frac{\kappa}{k}$ satisfies $|\rho_s| < 1$ given precisely the restriction $\rho_s \in [0,1)$ maintained in the text. Under $K > 0$, $\boldsymbol{e}^* = \boldsymbol{K}^{-1}\boldsymbol{b}$ is the unique global maximizer, validating Lemma 3 and the allocation in equation (14).

Consumption margin (Proposition 6). With investing and present consumption as competing tasks, the relevant cost Hessian (in the symmetric primitives $B = A$, $k_c = k$) has determinant

$$D = \delta k^2 - \kappa^2 > 0, \text{ equivalently } \rho_s < \sqrt{\delta}.$$

This is the positive-definiteness/SOC condition for the joint invest–consume problem and the maintained assumption of Proposition 6. It guarantees e^* and c^* in the proof of Proposition 6 are a joint maximum rather than a saddle point. Note $D > 0$ is strictly stronger than the single-task $\rho_s < 1$, reflecting that a stronger pull to consume (lower δ) tightens the concavity requirement.

A.4 Realization thresholds (Propositions 8–9)

The lock-in results are not interior optima of a smooth objective; they are indifference thresholds, so the relevant conditions are uniqueness of the switching hurdle and the signs of its comparative statics rather than a second-order condition.

Uniqueness of the hurdle (Prop. 8). Define $\Delta(r_H) = W_{switch(r_H)} - W_{hold}$. Both terminal-wealth expressions are affine and increasing in r_H, with W_{switch} having the larger slope whenever the investor retains positive reinvested capital; hence Δ is strictly monotone in r_H and crosses zero exactly once, at

$$r_H^* = \frac{r_L}{1 - \tau\theta_G}$$

Investors with opportunity above r_H^* switch and those below hold, so the population sorts on a single threshold. The comparative statics are

$$\frac{dr_H^*}{d\tau} = \frac{r_L\theta_G}{(1 - \tau\theta_G)^2} > 0, \quad \frac{dr_H^*}{d\theta_G} = \frac{r_L\tau}{(1 - \tau\theta_G)^2} > 0$$

confirming that a higher rate and a larger embedded gain both raise the hurdle (Prop. 8).

Deferral and step-up (Prop. 9). With escape probability ξ and effective holding tax $\hat{\tau} = (1 - \xi)\tau$, the same single-crossing argument applied to $W_{switch,don't-sell} = W_{never-sell}$ yields the unique hurdle

$$r_H^{\dagger} = r_H^{*} + \frac{\xi\tau\theta_G}{[(1 - \tau\theta_G)(1 - \hat{\tau})]}$$

which is strictly increasing in $\xi, \left(\frac{dr_H^{\dagger}}{d\xi} > 0\right)$. Deferral-capable investors therefore face a strictly higher switching threshold, as stated in Proposition 9.

A.5 Interior, zero, and subsidy regimes

The propositions report interior optimal rates; these are the relevant solution only where the rate is feasible ($\tau \in [0,1)$). The optimal policy partitions the parameter space into three regimes, with the boundaries following directly from the closed forms.

Taking the spillover-augmented welfare rate as the leading case, $\tau * = \frac{\lambda - 1 - \lambda\varphi}{2\lambda - 1}$ (with $s = 0$), the regime boundaries are:

- $\text{Interior tax } (0 < \tau^{*} < \frac{1}{2})$: holds for $0 \leq \phi < \frac{\lambda-1}{\lambda}$.
- $\text{Zero rate } (\tau^{*} = 0)$: at the threshold $\phi = \frac{\lambda-1}{\lambda}$, equivalently at $\lambda = 1$ when $\phi = 0$.
- $\text{Subsidy } (\tau^{*} < 0)$: for $\phi > \frac{\lambda-1}{\lambda}$ (Corollary 4.1). Under isoelastic cost the threshold generalizes to $\phi^{\dagger} = \frac{\lambda-1}{\lambda\nu}$ (Corollary 4.2), decreasing in ν.

The revenue-maximizing rate obeys the analogous partition with boundary $\phi = 1/\nu$ (so $\tau_{\phi}^{R} = \frac{1 - \nu\phi}{1+\nu}$ turns negative for $\phi > \frac{1}{\nu}$). These boundaries make explicit that the negative entries in Tables 2–3 are the model's intended subsidy region, not violations of an interior assumption: where the fiscal spillover or growth dividend is large enough, the optimal contract credits investment, and the interior-tax formulas are correctly superseded by the corner/subsidy solution.