Asset Returns, Portfolio Choice, and Proportional Wealth Taxation

Anders G Frøseth*

Draft — 2 March 2026

Abstract

We analyse the effect of a proportional wealth tax on asset returns, portfolio choice, and asset pricing in a partial equilibrium setting. The tax is levied annually on the market value of all holdings at a uniform rate. We show that such a tax is economically equivalent to the government acquiring a proportional stake in the investor's portfolio each period—a form of risk sharing in which expected wealth and risk are reduced by the same factor, while the return per share is unaffected. This multiplicative separability between the tax factor and the return realisation drives four main results. First, the coefficient of variation of wealth is invariant to the tax rate, since the tax reduces expected wealth and risk by the same proportion. Second, the optimal portfolio weights—and in particular the tangency portfolio—are independent of the tax rate. Third, the wealth tax is orthogonal to portfolio choice: in discrete time it induces a homothetic contraction of the opportunity set in the mean–standard deviation plane that preserves the Sharpe ratio of every portfolio. Fourth, both taxed and untaxed investors are willing to pay the same price per share for any asset. The results are derived first under geometric Brownian motion and then generalised to any return distribution in the location-scale family. A complementary Modigliani-Miller analysis, treating the tax claim as a separate security, confirms pricing neutrality and identifies an inconsistency in the existing literature regarding the discount rate used for after-tax cash flows. Imposing the Capital Asset Pricing Model as a special case confirms that after-tax betas equal pre-tax betas and the security market line contracts uniformly by $(1-\tau_w)$; under CRRA preferences, general-equilibrium returns and prices are unchanged. This resolves an error in Fama (2021), who overstates the price effect by adding the wealth tax to the cost of capital without adjusting the discount rate. The neutrality results depend on two conditions that are commonly violated in practice: universal taxation at market value, and frictionless markets. We formalise three channels through which relaxing these conditions breaks neutrality—book-value taxation, liquidity frictions, and dividend extraction—and show that they have opposing effects on asset prices.

JEL Classification: G11, G12, H21, H24.

Keywords: Wealth tax, portfolio choice, asset pricing, CAPM, Sharpe ratio, Modigliani–Miller, tax neutrality, book-value taxation.

*Independent Researcher. E-mail: indrefjorden@pm.me.

Contents

1 Introduction

This paper develops a framework for analysing the effect of a proportional wealth tax on asset returns, portfolio choice, and asset pricing. We compare two investors who are identical in all respects except that one (Investor A) pays an annual proportional wealth tax on the market value of all holdings, while the other (Investor B) does not.

The main results are:

1. A proportional wealth tax on all assets is economically equivalent to the government acquiring a proportional claim on the investor's assets, sharing both risk and return pro rata. From the investor's cash-flow perspective, this amounts to a periodic partial sale of shares to meet the tax obligation.

2. The tax does not alter the return distribution per unit of asset, the risk-reward profile (coefficient of variation), or the optimal portfolio weights.

3. The wealth tax is orthogonal to portfolio choice: it operates purely in the return dimension while portfolio optimisation operates in the risk dimension.

4. Both investors should be willing to pay the same price per share for any asset.

These results hold under the economic assumptions of proportional taxation on all assets and partial equilibrium. We develop them in two stages. In Sections 3–6, we derive the results under geometric Brownian motion (GBM), which provides a concrete and tractable setting with closed-form expressions for the moments of the wealth distribution. In Section 7, we show that the GBM assumption can be substantially relaxed: the results hold for any return distribution in the location-scale family (and for some results, any distribution with finite second moments). The distributional assumption turns out to play no substantive economic role—the structural engine behind all four results is the multiplicative separability of the wealth tax from the return realisation.

Section 8 discusses the conditions under which the results hold and clarifies their independence from any specific asset pricing model. A complementary analysis under the CAPM (Section 8.3) confirms that after-tax betas and the security market line are preserved, that general equilibrium is undisturbed under CRRA preferences, and identifies a pricing error in Fama (2021). A Modigliani-Miller perspective completes the discussion; a survey of the related literature is collected in Appendix D. Section 9 then asks what happens when the key conditions fail. Three channels of non-neutrality are formalised: book-value taxation (where the tax base diverges from market value), liquidity frictions (where forced selling incurs transaction costs), and dividend extraction (where the tax forces payouts that displace profitable investment). These channels have opposing effects on asset prices, and their net impact depends on the asset class. Section 10 summarises the contributions and outlines directions for further theoretical and empirical work. Appendix A analyses alternative tax payment mechanisms.

2 Setup

2.1 Asset Return Dynamics

There are K risky assets with prices following a multivariate GBM:

$$\frac{dP_i}{P_i} = \mu_i \, dt + \sum_j \sigma_{ij} \, dZ_j, \qquad i = 1, \ldots, K. \tag{1}$$

In vector notation:

$$\frac{d\mathbf{P}}{\mathbf{P}} = \boldsymbol{\mu} \, dt + \boldsymbol{\Sigma} \, d\mathbf{Z} \tag{2}$$

where $\boldsymbol{\mu}$ is the $K \times 1$ vector of expected returns and $\mathbf{V} = \boldsymbol{\Sigma}\boldsymbol{\Sigma}^\top$ is the $K \times K$ covariance matrix. A risk-free asset with continuous return r_f is also available.

The GBM assumption will be relaxed in Section 7, where we show that the main results depend only on the existence of well-defined first and second moments. In the multivariate setting, the natural distributional class is the family of elliptical distributions (Owen and Rabinovitch, 1983), which nests the multivariate normal (and hence GBM) as a special case.

2.2 Tax Structure

We consider a single tax instrument: a personal wealth tax at rate $\tau_w \in (0, 1)$, levied annually on the market value of all assets held by the investor (including the risk-free asset). There is no income tax and no capital gains tax. Two features of this specification are worth highlighting. First, the tax is *proportional*: the rate τ_w is the same regardless of wealth level. Second, the tax is *universal*: all assets are taxed at the same rate on market value, so no asset bears a higher or lower effective tax burden than any other. Together, proportionality and universality generate the multiplicative separability that drives the main results; relaxing either opens channels of non-neutrality (Section 9).

For expositional clarity, dividends are assumed to be paid out at year-end and consumed; the wealth tax is then paid by selling shares. This is a simplifying convention, not a necessary condition for neutrality: if dividends are instead reinvested in assets that remain within the tax base, the multiplicative separability is preserved and all results carry through (Appendix A).

2.3 Investors

Investor A pays wealth tax $\tau_w > 0$ on all asset holdings. **Investor B** pays no wealth tax ($\tau_w = 0$). The two investors are otherwise identical in preferences, information, and endowments.

2.4 Notation

Table 1 collects the principal symbols used throughout the paper.

Table 1: Notation guide.

Assets and returns		*Portfolio*	
K	Number of risky assets	$\mathbf{w}$	Portfolio weight vector
P_i, P_t	Asset price	$\mathbf{w}^*$	Optimal weights
$\boldsymbol{\mu}$, μ_i	Expected return (vector / scalar)	$\mathbf{w}_T$	Tangency portfolio
$\boldsymbol{\Sigma}$	Diffusion (volatility) matrix	$\mu_P(\mathbf{w})$	Portfolio expected return
$\mathbf{V} = \boldsymbol{\Sigma}\boldsymbol{\Sigma}^\top$	Covariance matrix	$\sigma_P(\mathbf{w})$	Portfolio volatility
r_f	Risk-free rate	$R_P(\mathbf{w})$	Portfolio rate of return (net)
$\mathbf{R}$	Return vector $(R_1, \ldots, R_K)^\top$	G_P	Portfolio gross return $1 + R_P$
$\mathbf{Z}$	Standard Brownian motion / standardised random vector		
$G^{(n)}$	Cumulative gross return P_n/P_0		
Wealth, tax, and investors		*After-tax quantities*	
W_0, W_n	Wealth at time 0, n	$R_W(\mathbf{w}, \tau_w)$	After-tax return on wealth
W_n^A, W_n^B	Wealth of Investor A, B	$\mu_W(\mathbf{w}, \tau_w)$	After-tax expected return
N_0, N_n	Number of shares at time 0, n	$\sigma_W(\mathbf{w}, \tau_w)$	After-tax volatility
τ_w	Wealth tax rate	r_f^A	After-tax risk-free rate
γ	Risk-aversion parameter	$\mathrm{SR}_W(\mathbf{w})$	After-tax Sharpe ratio
Statistical measures		*MM perspective (Section 8.4)*	
$E[\cdot]$	Expectation	V^0, V	Firm value (pre-/post-tax)
$\mathrm{SD}(\cdot)$	Standard deviation	$\bar{x}$, x_t	Expected / realised cash flow
$\mathrm{CV}(\cdot)$	Coefficient of variation	k, k^A	Cost of capital (pre-/after-tax)
$\mathrm{Cov}(\cdot)$	Covariance	β_U	Unlevered (asset) beta
$\mathrm{SR}(\cdot)$	Sharpe ratio	β_{tax}	Beta of the tax claim
CAPM (Section 8.3)		*Beyond neutrality (Section 9)*	
β_j, β_j^A	Asset beta; after-tax beta ($= \beta_j$)	B, θ	Book value; book-to-market B/V
MRP	Market risk premium $E[R_M] - r_f$	V_{PE}, V_{GE}	PE / GE asset value
MRP^A	After-tax MRP: $(1 - \tau_w)\cdot\mathrm{MRP}$	c_j	Illiquidity cost (stochastic)
β_j^τ	After-tax beta (heterog. tax)	δ	Payout ratio
τ_j	Asset-specific effective tax rate	ρ	Internal rate of return

Subscript P denotes portfolio (pre-tax) quantities; subscript W denotes after-tax wealth quantities. Superscripts A and B identify the taxed and untaxed investor, respectively. Bold symbols ($\mathbf{w}$, $\mathbf{R}$, $\boldsymbol{\mu}$, $\mathbf{V}$, $\boldsymbol{\Sigma}$, $\mathbf{Z}$) denote vectors or matrices; unbolded variants denote scalars.

3 Single Asset Case Under GBM

3.1 Wealth Dynamics

Consider a single risky asset with return parameters (μ, σ). The investor holds N_0 shares at date 0 with initial price P_0, so initial wealth is $W_0 = N_0 P_0$.

At each year-end $i = 1, 2, \ldots, n$:

1. The share price evolves to P_i according to the GBM.
2. Dividends are paid and consumed (simplifying convention; see Appendix A).
3. The wealth tax $\tau_w N_{i-1} P_i$ is due on the market value of holdings.
4. The investor sells a fraction τ_w of shares to pay the tax.

After n years, the number of shares held is:

$$N_n = N_0(1 - \tau_w)^n. \tag{3}$$

The investor's equity wealth is:

$$W_n = N_n P_n = W_0(1 - \tau_w)^n \frac{P_n}{P_0}. \tag{4}$$

Since P_n/P_0 is determined entirely by the GBM and is independent of the investor's tax status, we can write:

$$W_n = W_0(1 - \tau_w)^n \exp\left[\left(\mu - \tfrac{1}{2}\sigma^2\right) n + \sigma Z_n\right] \tag{5}$$

where $Z_n \sim N(0, n)$.

3.2 Moments of the Wealth Distribution

Expected wealth:

$$E[W_n] = W_0(1 - \tau_w)^n \, e^{\mu n}. \tag{6}$$

Standard deviation of wealth:

$$\mathrm{SD}(W_n) = W_0(1 - \tau_w)^n \, e^{\mu n} \sqrt{e^{\sigma^2 n} - 1}. \tag{7}$$

Coefficient of variation:

$$\mathrm{CV}(W_n) = \frac{\mathrm{SD}(W_n)}{E[W_n]} = \sqrt{e^{\sigma^2 n} - 1}. \tag{8}$$

3.3 Key Result: Scaling Invariance

The wealth tax enters only through the deterministic prefactor $(1 - \tau_w)^n$, which multiplies both $E[W_n]$ and $\mathrm{SD}(W_n)$ identically. Therefore:

Proposition 1 (CV Invariance, GBM)**.** *Under GBM with a proportional wealth tax on all assets, the coefficient of variation of wealth at any horizon is invariant to the wealth tax rate:*

$$\mathrm{CV}_A(W_n) = \mathrm{CV}_B(W_n) = \sqrt{e^{\sigma^2 n} - 1}. \tag{9}$$

The wealth tax reduces absolute expected wealth and absolute risk (standard deviation) by the same proportion $(1 - \tau_w)^n$, leaving the relative risk-reward profile unchanged.

This result generalises to any return distribution with finite second moments; see Proposition 5 in Section 7.4.

3.4 Economic Interpretation: Proportional Dilution and Risk Sharing

The wealth tax is economically equivalent to an annual transfer of a fraction τ_w of the investor's position to the government. Each share retained has the same return distribution as it would in the absence of the tax. The tax operates on the *quantity* of the position, not on the *quality* of the return.

This equivalence is exact under the following conditions:

- The tax rate τ_w is proportional (no exemptions, no progressive rates, no caps).
- The tax applies to all assets at the same rate.
- The asset price process P_t is unaffected by the investor's tax status (partial equilibrium).

Remark (Risk sharing). The proportional-dilution mechanism admits a dual interpretation. From the Modigliani-Miller perspective developed in Section 8.4, the same mechanism can be viewed as the government holding a proportional claim on the investor's assets. The government participates in both upside and downside pro rata: the tax reduces the investor's expected return and risk exposure by the same factor $(1 - \tau_w)$. The wealth tax is therefore a form of risk sharing between the investor and the state, not a one-sided extraction. These two descriptions—proportional dilution (the investor's cash-flow perspective) and proportional claim (the valuation perspective)—are economically equivalent.

Figure 1 illustrates the mechanism over several periods.

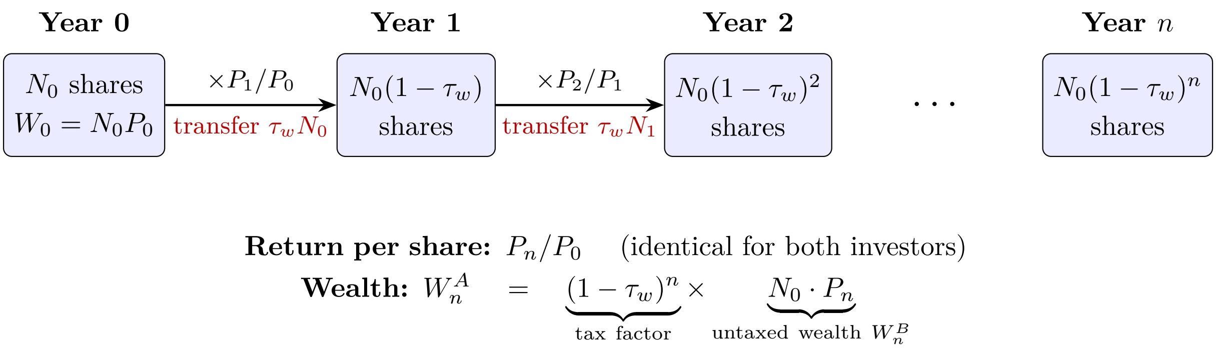

Figure 1: The proportional-dilution mechanism. At each period end, a fraction τ_w of the investor's shares is transferred to the government as tax. The share price process P_t is unaffected. After n periods, the number of shares has decayed by the deterministic factor $(1 - \tau_w)^n$, while the return per share remains P_n/P_0—identical for taxed and untaxed investors.

4 Multi-Asset Portfolio Choice Under GBM

4.1 Wealth Dynamics with Portfolio Weights

Investor B holds portfolio weights $\mathbf{w}$ in risky assets and $(1-\mathbf{1}^\top\mathbf{w})$ in the risk-free asset. Following the continuous-time framework of Merton (1969), wealth evolves as:

$$\frac{dW_B}{W_B} = [r_f + \mathbf{w}^\top(\boldsymbol{\mu} - r_f\mathbf{1})]\,dt + \mathbf{w}^\top\boldsymbol{\Sigma}\,d\mathbf{Z}. \tag{10}$$

For Investor A, the wealth tax applies to total wealth across all assets at the same rate τ_w. In the continuous-time approximation:

$$\frac{dW_A}{W_A} = [r_f + \mathbf{w}^\top(\boldsymbol{\mu} - r_f\mathbf{1}) - \tau_w]\,dt + \mathbf{w}^\top\boldsymbol{\Sigma}\,d\mathbf{Z}. \tag{11}$$

Define the portfolio expected return and volatility:

$$\mu_P(\mathbf{w}) = r_f + \mathbf{w}^\top(\boldsymbol{\mu} - r_f\mathbf{1}), \tag{12}$$

$$\sigma_P(\mathbf{w}) = \sqrt{\mathbf{w}^\top\mathbf{V}\mathbf{w}}. \tag{13}$$

Then the wealth dynamics decompose as:

$$\mu_W(\mathbf{w}, \tau_w) = \mu_P(\mathbf{w}) - \tau_w, \tag{14}$$

$$\sigma_W(\mathbf{w}, \tau_w) = \sigma_P(\mathbf{w}). \tag{15}$$

The wealth tax enters the drift additively and does not appear in the diffusion coefficient.

4.2 Optimality Conditions

The investor chooses $\mathbf{w}$ to maximise an objective $f(\mu_W, \sigma_W)$ that is increasing in μ_W and decreasing in σ_W (Markowitz, 1952). For concreteness, consider the mean-variance objective with risk aversion parameter γ:

$$\max_{\mathbf{w}} \quad \mu_P(\mathbf{w}) - \tau_w - \frac{\gamma}{2}\sigma_P^2(\mathbf{w}). \tag{16}$$

The first-order condition is:

$$\nabla_{\mathbf{w}}\mu_P = \gamma\,\mathbf{V}\mathbf{w} \qquad \Longrightarrow \qquad \boldsymbol{\mu} - r_f\mathbf{1} = \gamma\,\mathbf{V}\mathbf{w}^*. \tag{17}$$

Solving:

$$\mathbf{w}^* = \frac{1}{\gamma}\mathbf{V}^{-1}(\boldsymbol{\mu} - r_f\mathbf{1}). \tag{18}$$

Proposition 2 (Portfolio Invariance, GBM)**.** *The optimal portfolio weights $\mathbf{w}^*$ are independent of the wealth tax rate τ_w:*

$$\mathbf{w}_A^* = \mathbf{w}_B^* = \frac{1}{\gamma}\mathbf{V}^{-1}(\boldsymbol{\mu} - r_f\mathbf{1}). \tag{19}$$

This holds for any objective function $f(\mu_W, \sigma_W)$ where τ_w enters only through an additive shift to μ_W.

This full weight invariance is specific to the continuous-time formulation. In discrete time, the tangency portfolio is still invariant, but the total risky allocation may depend on preferences; see Proposition 6 in Section 7.5 and the comparison in Appendix C.

4.3 Orthogonality

The invariance of $\mathbf{w}^*$ follows from a deeper structural property. The gradients with respect to $\mathbf{w}$ are:

$$\nabla_{\mathbf{w}}\mu_W = \boldsymbol{\mu} - r_f\mathbf{1}, \qquad \nabla_{\mathbf{w}}\sigma_W = \frac{\mathbf{V}\mathbf{w}}{\sigma_P}. \tag{20}$$

Neither depends on τ_w. The wealth tax shifts the objective in the μ-direction by a constant $-\tau_w$, which is orthogonal to the control variable $\mathbf{w}$ that operates through both μ_P and σ_P.

Proposition 3 (Orthogonality, GBM)**.** *In the (σ, μ) plane, the wealth tax is a vertical translation of the entire opportunity set (efficient frontier, capital allocation line, and risk-free rate) by $-\tau_w$. The indifference curves of any risk-averse investor translate by the same amount. The tangency point—and hence the optimal portfolio—moves purely vertically. The tax is orthogonal to portfolio choice.*

Formally, define the direction of the wealth tax effect as $\mathbf{e}_\mu = (0, -1)$ in (σ, μ) space. Portfolio choice operates along the efficient frontier, whose tangent direction has a nonzero σ-component at any interior point. The inner product $\mathbf{e}_\mu \cdot \nabla_{\mathbf{w}}\sigma_P = 0$, confirming orthogonality.

Figure 2 illustrates this geometry.

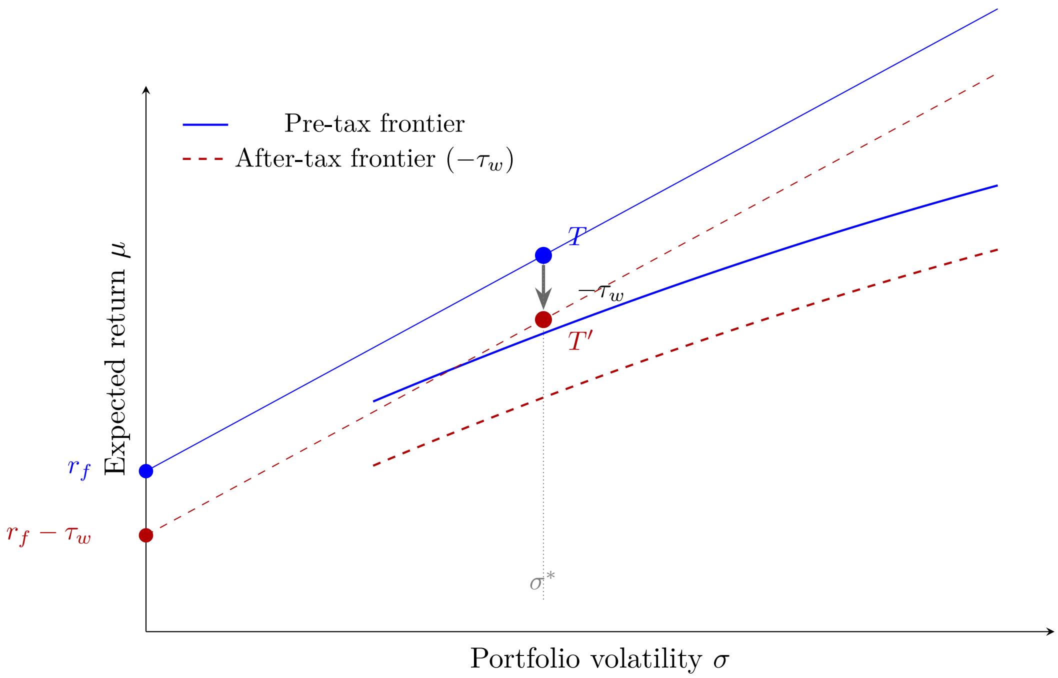

Figure 2: Orthogonality in continuous time. The wealth tax translates the entire opportunity set vertically by $-\tau_w$: the efficient frontier, the risk-free rate, and the capital allocation line all shift down by the same amount. The tangency portfolio T moves to T' at the same volatility σ^*—the tax is orthogonal to portfolio choice.

In discrete time, the geometric interpretation changes from a vertical translation to a homothetic contraction; see Proposition 7 in Section 7.6 and the comparison in Appendix C.

5 Asset Pricing Implications Under GBM

5.1 Per-Share Valuation

Consider one share purchased at price P_0. After one year:

- **Investor B** holds 1 share worth P_1. Return: P_1/P_0.
- **Investor A** holds 1 share worth P_1, pays tax $\tau_w P_1$, retains $(1-\tau_w)$ shares. Return per share held: still P_1/P_0.

Investor A receives this return on a shrinking number of shares, but the return *per share*—and hence the price they are willing to pay for any individual share—is unchanged.

Proposition 4 (Pricing Neutrality, GBM)**.** *Under GBM with a proportional wealth tax on all assets, both investors are willing to pay the same price per share for any asset. The wealth tax does not create a pricing wedge.*

This result—which parallels the DCF-based finding of Bjerksund and Schjelderup (2022)—is in fact completely distribution-free; see Proposition 8 in Section 7.7.

5.2 Discussion

This result may appear counterintuitive. The common argument is: "Investor A faces an additional cost (the wealth tax), so they should demand a higher return, which implies they would pay a lower price." The error in this reasoning is the conflation of the return on the *asset* with the return on *wealth*:

- The **return on the asset** is P_1/P_0, which is identical for both investors.
- The **return on wealth** is P_1/P_0 for Investor B and $(1-\tau_w)P_1/P_0$ for Investor A.

The wealth tax reduces the return on wealth, not the return on the asset. Since asset prices reflect asset returns—not any particular investor's wealth accumulation rate—both investors value the asset identically.

The analogy is exact: the wealth tax is equivalent to any other proportional personal expense (such as a fixed consumption rate proportional to wealth). Such expenses reduce wealth accumulation but do not affect asset valuations.

6 Summary of GBM Results

7 Generalisation Beyond Geometric Brownian Motion

7.1 Motivation

The results in Sections 3–5 were derived under the assumption that asset prices follow a geometric Brownian motion. While GBM provides a tractable and well-understood framework, it is a

Table 2: Comparison of taxed and untaxed investors under GBM.

Property	Investor A (wealth tax τ_w)	Investor B (no tax)
Return per share	P_n/P_0	P_n/P_0
Shares held after n years	$N_0(1-\tau_w)^n$	N_0
Expected wealth $E[W_n]$	$W_0(1-\tau_w)^n e^{\mu n}$	$W_0 e^{\mu n}$
SD(W_n)	$W_0(1-\tau_w)^n e^{\mu n}\sqrt{e^{\sigma^2 n}-1}$	$W_0 e^{\mu n}\sqrt{e^{\sigma^2 n}-1}$
Coefficient of variation	$\sqrt{e^{\sigma^2 n}-1}$	$\sqrt{e^{\sigma^2 n}-1}$
Optimal weights $\mathbf{w}^*$	$\frac{1}{\gamma}\mathbf{V}^{-1}(\boldsymbol{\mu}-r_f\mathbf{1})$	$\frac{1}{\gamma}\mathbf{V}^{-1}(\boldsymbol{\mu}-r_f\mathbf{1})$
Price per share	P	P
Sharpe ratio	$\frac{\mu_T-r_f}{\sigma_T}$	$\frac{\mu_T-r_f}{\sigma_T}$

μ_T and σ_T denote the tangent portfolio return and volatility.

strong assumption: it restricts returns to be lognormally distributed, rules out fat tails, and imposes a continuous-time diffusion structure on prices.

We now show that none of these features are essential to the main results. The four propositions depend not on the specific distributional form of returns, but on the algebraic structure of how a proportional wealth tax interacts with the return process. Specifically, the results rest on two pillars:

1. **The wealth tax is a multiplicative scalar on wealth**, deterministic and independent of the return realisation.
2. **Asset returns have well-defined first and second moments**, so that the mean-variance framework applies.

The first pillar is a consequence of proportionality and universality—the same conditions already identified in the GBM analysis. The second replaces the GBM assumption with a much weaker requirement that encompasses any distribution in the **location-scale family** (normal, Student-t, logistic, uniform, and more generally, the class of elliptical distributions).

We re-derive each proposition, clearly identifying the minimal assumptions required. It will emerge that Propositions 1′ and 4′ require essentially no distributional assumption at all, while Propositions 2′ and 3′ exploit the location-scale structure to justify working in the (σ, μ) plane.

7.2 Generalised Assumptions

We retain the economic assumptions of Section 2 (proportional taxation, universality, partial equilibrium, dividends consumed) and replace the GBM return dynamics with the following:

Assumption A1 (Finite moments). The vector of asset returns $\mathbf{R} = (R_1, \ldots, R_K)^\top$ has a joint distribution with well-defined mean vector $E[\mathbf{R}] = \boldsymbol{\mu}$ and positive-definite covariance matrix $\mathrm{Cov}(\mathbf{R}) = \mathbf{V}$. No further assumption is made about the form of the distribution.

Assumption A2 (Location-scale family). The return distribution belongs to the location-scale family: there exists a standardised random vector $\mathbf{Z}$ with $E[\mathbf{Z}] = \mathbf{0}$ and $\mathrm{Cov}(\mathbf{Z}) = \mathbf{I}$ such that

$$\mathbf{R} \stackrel{d}{=} \boldsymbol{\mu} + \boldsymbol{\Sigma}\mathbf{Z} \tag{21}$$

where $\boldsymbol{\Sigma}$ is a matrix with $\boldsymbol{\Sigma}\boldsymbol{\Sigma}^\top = \mathbf{V}$, and $\mathbf{Z}$ may follow any distribution with well-defined second moments. The shape of the distribution (e.g. tail behaviour, kurtosis) is encoded in $\mathbf{Z}$ and is invariant under affine transformations.

Remark. Assumption A2 is a strengthening of A1 that is needed only for Propositions $2'$ and $3'$. When A2 holds, the distribution is fully characterised (up to shape) by its first two moments, which justifies the use of mean-variance analysis (Meyer, 1987; Chamberlain, 1983). In the multivariate setting, the natural extension is the class of **elliptical distributions**, which nests the multivariate normal (and hence GBM) as a special case (Owen and Rabinovitch, 1983; Hamada and Valdez, 2008). The precise role of the distributional assumption is discussed in Appendix B.

Assumption A3 (Tax structure). Unchanged from Section 2.2: a proportional wealth tax at rate $\tau_w \in (0,1)$ is levied on the market value of all assets (including the risk-free asset) at each period end. The tax is paid by selling a fraction τ_w of all positions.

Assumption A4 (Partial equilibrium). The return distribution $(\boldsymbol{\mu}, \mathbf{V})$ is exogenous—it does not depend on the tax rate.

7.3 The Tax as a Multiplicative Scalar

The proportional-dilution interpretation from Section 3.4 is entirely distribution-free. At each period end, a fraction τ_w of all positions is transferred to the government as tax. After n periods, the number of shares held is:

$$N_n = N_0(1-\tau_w)^n \tag{22}$$

and wealth is:

$$W_n^A = N_0(1-\tau_w)^n P_n = (1-\tau_w)^n \cdot W_0 \cdot \frac{P_n}{P_0}. \tag{23}$$

Define the **cumulative gross return** $G^{(n)} \equiv P_n/P_0$. This is a random variable whose distribution is determined by the asset return process and is, by Assumption A4, independent of the investor's tax status. Then:

$$W_n^A = (1-\tau_w)^n \cdot W_0 \cdot G^{(n)} = (1-\tau_w)^n \cdot W_n^B \tag{24}$$

where $W_n^B = W_0 \cdot G^{(n)}$ is the wealth of the untaxed investor. The wealth tax enters as a **deterministic multiplicative scalar** $(1-\tau_w)^n$ that is independent of the return realisation $G^{(n)}$.

This multiplicative separability is the structural engine behind all four results. It requires only proportionality (the tax rate is constant) and universality (all assets are taxed at the same rate). It does not require GBM, continuity, or any specific distributional form.

Under GBM, the cumulative gross return takes the specific form $G^{(n)} = \exp[(\mu - \frac{1}{2}\sigma^2)n + \sigma Z_n]$, but the results below hold for any $G^{(n)}$ with well-defined moments.

7.4 Generalised Proposition 1: CV Invariance

Requires: A1 (finite moments), A3 (proportional tax). No distributional assumption.

Since $W_n^A = (1-\tau_w)^n \cdot W_n^B$ and $(1-\tau_w)^n > 0$ is a deterministic scalar:

$$E[W_n^A] = (1-\tau_w)^n \, E[W_n^B], \tag{25}$$

$$\mathrm{SD}(W_n^A) = (1-\tau_w)^n \, \mathrm{SD}(W_n^B), \tag{26}$$

$$\mathrm{CV}(W_n^A) = \frac{\mathrm{SD}(W_n^A)}{E[W_n^A]} = \frac{\mathrm{SD}(W_n^B)}{E[W_n^B]} = \mathrm{CV}(W_n^B). \tag{27}$$

Proposition 5 (Generalised CV Invariance)**.** *Let asset returns have any joint distribution with well-defined first and second moments (Assumption A1). Under a proportional wealth tax on all assets (Assumption A3), the coefficient of variation of wealth at any horizon n is invariant to the tax rate:*

$$\mathrm{CV}_A(W_n) = \mathrm{CV}_B(W_n) = \frac{\mathrm{SD}(G^{(n)})}{E[G^{(n)}]} \tag{28}$$

where $G^{(n)} = P_n/P_0$ is the cumulative gross return. The result holds for any return distribution—normal, lognormal, fat-tailed, skewed, discrete, or continuous—provided the first two moments exist.

Proof. The proof is a single line of algebra exploiting the linearity of the expectation operator and the absolute homogeneity of the standard deviation:

$$\mathrm{CV}(cX) = \frac{|c|\,\mathrm{SD}(X)}{c\,E[X]} = \frac{\mathrm{SD}(X)}{E[X]} = \mathrm{CV}(X) \tag{29}$$

for any positive constant c and random variable X with $E[X] > 0$. Setting $c = (1-\tau_w)^n$ and $X = W_n^B$ completes the argument. □

Remark. Under GBM, Proposition 1 gave the specific formula $\mathrm{CV}(W_n) = \sqrt{e^{\sigma^2 n} - 1}$, which depended on the volatility parameter σ and the horizon n through the lognormal moment structure. Proposition 5 shows that the *invariance to* τ_w does not depend on this formula—it is a consequence of the multiplicative structure of the tax alone.

7.5 Generalised Proposition 2: Portfolio Weight Invariance

Requires: A1 (finite moments), A3 (proportional tax). Full result requires A2 (location-scale) or CRRA preferences.

7.5.1 Setup

Consider K risky assets and a risk-free asset with one-period gross return $(1+r_f)$. The investor allocates portfolio weights $\mathbf{w} = (w_1, \ldots, w_K)^\top$ to risky assets, with the remainder $(1 - \mathbf{1}^\top \mathbf{w})$ in the risk-free asset. The portfolio rate of return is:

$$R_P(\mathbf{w}) = r_f + \mathbf{w}^\top(\mathbf{R} - r_f\mathbf{1}) \tag{30}$$

with mean and variance:

$$\mu_P(\mathbf{w}) = r_f + \mathbf{w}^\top(\boldsymbol{\mu} - r_f\mathbf{1}), \qquad \sigma_P^2(\mathbf{w}) = \mathbf{w}^\top\mathbf{V}\mathbf{w}. \tag{31}$$

This notation is consistent with the portfolio return and volatility defined in Section 4.1, but now $\mathbf{R}$ may follow any distribution satisfying A1.

7.5.2 After-Tax Returns

After paying the wealth tax, end-of-period wealth is:

$$W_1^A = (1 - \tau_w) \cdot W_0(1 + R_P) = (1 - \tau_w) \cdot W_1^B. \tag{32}$$

The after-tax rate of return (on beginning-of-period wealth) is:

$$R_W(\mathbf{w}, \tau_w) = (1 - \tau_w)(1 + R_P(\mathbf{w})) - 1 \tag{33}$$

with moments:

$$\mu_W(\mathbf{w}, \tau_w) = (1 - \tau_w)(1 + \mu_P(\mathbf{w})) - 1, \tag{34}$$

$$\sigma_W(\mathbf{w}, \tau_w) = (1 - \tau_w)\,\sigma_P(\mathbf{w}). \tag{35}$$

The after-tax risk-free return (setting $\mathbf{w} = \mathbf{0}$) is:

$$r_f^A = (1 - \tau_w)(1 + r_f) - 1 = r_f - \tau_w(1 + r_f). \tag{36}$$

Note the contrast with the continuous-time GBM case (Section 4.1), where the tax entered the drift additively ($\mu_W = \mu_P - \tau_w$) and did not affect volatility ($\sigma_W = \sigma_P$). In the discrete-time formulation, the tax scales both the mean and the standard deviation by $(1 - \tau_w)$. This distinction is important for the portfolio choice result but does not affect the Sharpe ratio, as we now show.

7.5.3 Sharpe Ratio Preservation

The after-tax excess return of portfolio $\mathbf{w}$ over the after-tax risk-free rate is:

$$\mu_W - r_f^A = (1 - \tau_w)(1 + \mu_P) - 1 - [r_f - \tau_w(1 + r_f)] = (1 - \tau_w)(\mu_P - r_f). \tag{37}$$

The after-tax Sharpe ratio (Sharpe, 1966) is therefore:

$$\mathrm{SR}_W(\mathbf{w}) = \frac{\mu_W - r_f^A}{\sigma_W} = \frac{(1 - \tau_w)(\mu_P - r_f)}{(1 - \tau_w)\sigma_P} = \frac{\mu_P - r_f}{\sigma_P} = \mathrm{SR}_P(\mathbf{w}). \tag{38}$$

The $(1-\tau_w)$ factors cancel exactly. The Sharpe ratio of every portfolio is invariant to the wealth tax.

7.5.4 Tangency Portfolio Invariance

The tangency portfolio $\mathbf{w}_T$ is the portfolio that maximises the Sharpe ratio:

$$\mathbf{w}_T = \arg\max_{\mathbf{w}} \frac{\mu_P(\mathbf{w}) - r_f}{\sigma_P(\mathbf{w})} = \arg\max_{\mathbf{w}} \mathrm{SR}_P(\mathbf{w}). \tag{39}$$

Since $\mathrm{SR}_W(\mathbf{w}) = \mathrm{SR}_P(\mathbf{w})$ for all $\mathbf{w}$, the tangency portfolio is the same in the after-tax space as in the pre-tax space. The standard solution is:

$$\mathbf{w}_T \propto \mathbf{V}^{-1}(\boldsymbol{\mu} - r_f\mathbf{1}) \tag{40}$$

which does not depend on τ_w.

7.5.5 Full Weight Invariance Under CRRA

For the allocation between the tangency portfolio and the risk-free asset, we need to specify preferences. Under CRRA utility $U(W) = W^{1-\gamma}/(1-\gamma)$:

$$E[U(W_1^A)] = \frac{[(1-\tau_w)W_0]^{1-\gamma}}{1-\gamma}\, E\Big[(1+R_P(\mathbf{w}))^{1-\gamma}\Big]\,. \tag{41}$$

The prefactor $[(1-\tau_w)W_0]^{1-\gamma}$ is a positive constant that does not affect the arg max over $\mathbf{w}$. For log utility ($\gamma = 1$):

$$E[\log W_1^A] = \log[(1-\tau_w)W_0] + E[\log(1+R_P(\mathbf{w}))]. \tag{42}$$

Again, the first term is constant. Therefore:

Proposition 6 (Generalised Portfolio Invariance)**.** *Let asset returns have any joint distribution with well-defined first and second moments (Assumption A1), and let the wealth tax be proportional on all assets (Assumption A3).*

(a) The Sharpe ratio of every portfolio is invariant to the wealth tax rate:

$$\mathrm{SR}_W(\mathbf{w}) = \mathrm{SR}_P(\mathbf{w}) \qquad \forall\, \mathbf{w}. \tag{43}$$

The tangency portfolio—the optimal allocation among risky assets—is identical for both investors.

(b) Under CRRA preferences (including log utility), the full optimal weight vector $\mathbf{w}^*$ *(including the allocation between risky and risk-free assets) is independent of* τ_w*:*

$$\mathbf{w}_A^* = \mathbf{w}_B^*. \tag{44}$$

These results hold without any specific distributional assumption on returns.

Remark (Mean-variance preferences). For a mean-variance investor maximising $\mu_W - \frac{\gamma}{2}\sigma_W^2$ over after-tax returns, the first-order condition yields $\mathbf{w}^* = \frac{1}{\gamma(1-\tau_w)}\mathbf{V}^{-1}(\boldsymbol{\mu} - r_f\mathbf{1})$. The tax affects the total risky allocation through the factor $1/(1-\tau_w)$, but the *composition* of the risky portfolio (the tangency portfolio) is unchanged. In the continuous-time limit, this scaling effect vanishes and the full weight vector is invariant, recovering Proposition 2. See Appendix C for a systematic comparison.

7.6 Generalised Proposition 3: Orthogonality

Requires: A2 (location-scale family), to justify the (σ, μ) representation.

7.6.1 The Discrete-Time Geometry

Under the location-scale assumption (A2), any investor's feasible set can be represented in the mean-standard deviation plane. The efficient frontier, the capital allocation line, and the indifference curves all live in this plane.

Consider the transformation induced by the wealth tax on after-tax returns. Writing in terms of the **gross return** $G_P = 1 + R_P$ and the after-tax gross return $G_W = (1-\tau_w)G_P$:

$$E[G_W] = (1-\tau_w)\, E[G_P], \qquad \mathrm{SD}(G_W) = (1-\tau_w)\,\mathrm{SD}(G_P). \tag{45}$$

In the (σ, G) plane, the wealth tax maps every point $(\sigma_P, \bar{G}_P)$ to $((1-\tau_w)\sigma_P,\ (1-\tau_w)\bar{G}_P)$. This is a **homothetic contraction** centred at the origin, with factor $(1-\tau_w) < 1$.

7.6.2 Properties of the Contraction

The homothetic contraction has several important geometric properties:

1. **Every point moves radially inward toward the origin.** The direction from the origin to any point is preserved; only the distance changes.
2. **The slope of every ray through the origin is preserved.** In particular, the slope of the capital allocation line (which is the Sharpe ratio plus a constant related to r_f) is invariant.
3. **The efficient frontier contracts uniformly.** Every point on the frontier maps to a point on the after-tax frontier at the same angular position relative to the origin.
4. **The tangency point is preserved.** Since both the frontier and the risk-free point contract along their respective rays, and the slope of the line connecting them is preserved, the tangency occurs at the same portfolio.

7.6.3 Comparison with the Continuous-Time Case

In continuous time (Section 4.3), the wealth tax induced a **pure vertical translation** of the entire (σ, μ) plane by $-\tau_w$: the efficient frontier shifted down, the risk-free rate shifted down, and the tangency point moved straight down. The tax direction $\mathbf{e}_\mu = (0, -1)$ was literally orthogonal to the frontier's tangent, which had a nonzero σ-component.

In discrete time, the transformation is richer: it is a contraction (scaling of both axes), not a translation (shift of one axis). Nevertheless, the economic content is the same:

Proposition 7 (Generalised Orthogonality)**.** *Under the location-scale assumption (A2), the wealth tax induces a homothetic contraction of the opportunity set in the $(\sigma, 1+\mu)$ plane by the factor $(1-\tau_w)$. The contraction preserves:*

- *the Sharpe ratio of every portfolio,*
- *the composition of the tangency portfolio,*
- *the slope of the capital allocation line.*

The tax operates purely as a radial scaling of the opportunity set. Portfolio optimisation selects the tangency point based on the angular position of the frontier relative to the risk-free ray—a quantity that is invariant under radial scaling. In this sense, the wealth tax remains orthogonal to portfolio choice.

In the continuous-time limit ($\tau_w\,dt \to 0$), the contraction degenerates into a vertical translation, recovering the exact orthogonality of Proposition 3. Figure 3 illustrates the discrete-time geometry; compare with the continuous-time case in Figure 2.

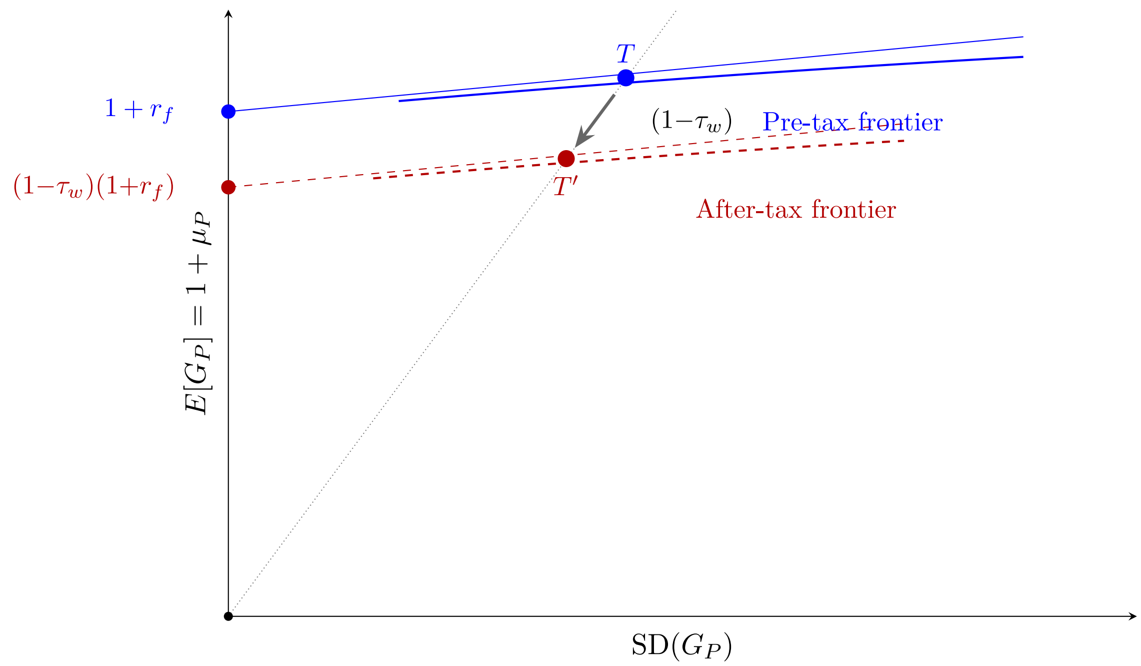

Figure 3: Orthogonality in discrete time. In the $(\sigma, 1+\mu)$ plane, the wealth tax induces a homothetic contraction of the opportunity set toward the origin by the factor $(1-\tau_w)$. Every point moves radially inward: the efficient frontier, the risk-free gross return, and the tangency portfolio T all contract along their respective rays from the origin. The slope of the capital allocation line—and hence the Sharpe ratio—is preserved.

7.7 Generalised Proposition 4: Pricing Neutrality

Requires: A3 (proportional tax), A4 (partial equilibrium). No distributional assumption.

The pricing result is the most general of the four. The argument of Section 5.1 is entirely distribution-free: the return per share is P_1/P_0 regardless of the investor's tax status, because the tax operates on the *number* of shares, not on the *return per share*. No assumption about the distribution of P_1 is needed.

For discounted cash flow valuation, this result is confirmed by Bjerksund and Schjelderup (2022), who show that the domestic (taxed) investor's NPV equals that of the foreign (untaxed) investor. The mechanism is discussed in Section 8.2.

Proposition 8 (Generalised Pricing Neutrality)**.** *Under a proportional wealth tax on all assets (Assumption A3) and exogenous asset return distributions (Assumption A4), both investors are willing to pay the same price per share for any asset. This result is completely distribution-free: it does not require GBM, finite moments, or any specific return process.*

Remark (Tax base timing)*.* In the model, the wealth tax is levied at period end on the current market value of holdings (Assumption A3). The neutrality result survives because the tax operates as a fixed fractional sale of shares: the investor sells a fraction τ_w of all positions regardless of the realised price, so the multiplicative structure $(1-\tau_w)^n$ is deterministic and separates cleanly from the stochastic return. Kruschwitz et al. (2023) show that using end-of-period market values can create arbitrage in a cost-of-capital framework that sets $k^\tau = k - \tau_w$. Our after-tax discount rate $k^A = (1-\tau_w)(1+k) - 1$ (equation (57)) differs from the naive $k - \tau_w$ by a term $-\tau_w k$, and is consistent with no-arbitrage under end-of-period assessment. See Section 8.1 for a fuller discussion and the connection to Norwegian institutional practice.

The role of distributional assumptions—which results require the location-scale family and which are distribution-free—is discussed in Appendix B.

8 Discussion

8.1 Essential Economic Conditions

The results in Sections 3–7 hold under specific economic assumptions. Each assumption, if relaxed, opens a channel through which the wealth tax *does* affect the risk-reward profile or portfolio choice:

Proportionality. The tax must be a fixed fraction of market value. Exemption thresholds, progressive rates, or caps introduce nonlinear transformations that change the shape of the wealth distribution and break CV invariance.

Universality. The tax must apply to all assets at the same rate. If some assets are exempt (e.g., the risk-free asset, retirement accounts, real estate), the tax creates differential costs across assets, which distorts relative Sharpe ratios and portfolio weights.

Tax base specification. The results assume the wealth tax base is the market value of holdings.

Two aspects of the tax base matter: the *valuation basis* (market value vs. book value, analysed in Section 9.2) and the *timing* (beginning vs. end of period).

Kruschwitz et al. (2023) show that in a cost-of-capital framework, using end-of-period market values with the naive discount rate $k^\tau = k - \tau_w$ creates arbitrage opportunities. Our model levies the tax at period end on current market values (Assumption A3), but avoids this problem for two reasons. First, the proportional-dilution mechanism—transferring a fixed fraction τ_w of shares regardless of the realised price—generates the deterministic factor $(1-\tau_w)^n$ that separates from the stochastic return without requiring a modified discount rate. Second, where a discount rate is needed (the MM perspective, Section 8.4), we derive $k^A = (1-\tau_w)(1+k) - 1$, which differs from $k - \tau_w$ by the correction term $-\tau_w k$ and is arbitrage-free under end-of-period assessment.

The alternative arbitrage-free discount rate derived by Kruschwitz et al. (2023) for end-of-period assessment, $k^K = k - \tau_w(1+k)/(1+r_f)$, rests on a narrower tax base. In their pricing equation, $(1-\tau_w)$ multiplies only the continuation market value V^τ_{t+1}, not the cash flow CF_{t+1}: the wealth tax is levied on the ex-dividend market value of the security, while the dividend escapes the tax base entirely. Their formula therefore requires that no part of the dividend remains in the investor's taxable estate at the assessment date—an implicit assumption that is difficult to reconcile with a comprehensive wealth tax on net worth. If dividends are instead retained (as cash, deposits, or reinvested capital) and therefore remain within the tax base, then $(1-\tau_w)$ applies to the entire gross payoff and our multiplicative formula obtains (see Appendix A for a detailed analysis of alternative payment mechanisms).

Remark (Consistency of the Kruschwitz et al. setup). The arbitrage identified by Kruschwitz et al. (2023) is in fact present in their model even at $\tau_w = 0$, before any tax is introduced. Their two-state binomial setup specifies three securities (one bond and two risky assets) with independently chosen costs of capital k_1 and k_2. In a two-state world, however, no-arbitrage constrains the risk premia:

$$\frac{k_1 - r_f}{1 + k_1} = \frac{u_1}{u_2} \cdot \frac{k_2 - r_f}{1 + k_2}, \tag{46}$$

where u_1/u_2 is the ratio of up-state payoffs. For the numerical parameters used by Kruschwitz et al., this condition fails. Setting their arbitrage-profit expression (their equation 9) to $\tau_w = 0$ yields a non-zero profit, confirming that the overcomplete market is already mispriced. The arbitrage they attribute to the naive discount rate $k^\tau = k - \tau_w$ is therefore an artefact of a pre-existing model inconsistency, not a consequence of wealth taxation. We thank Petter Bjerksund for drawing our attention to this observation.

Bjerksund and Schjelderup (2022) assume the tax base is the beginning-of-period market value, for which $k - \tau_w$ is the correct arbitrage-free discount rate (Kruschwitz et al., 2023). In the Norwegian institutional setting, both timing conventions coexist. Listed shares held *directly* are assessed at market value on 1 January of the assessment year—effectively the end-of-period price (*skatteloven* § 4-12(1), § 4-1). However, most substantial Norwegian investors hold listed equities through unlisted holding companies under the participation exemption (*fritaksmetoden*). The holding company's shares are assessed at tax book value per 1 January of the *income* year—the beginning of the period—so the effective tax base for these investors is a lagged, predetermined quantity, consistent with the Bjerksund and Schjelderup assumption. For unlisted operating companies, the tax base is the company's net asset value at 1 January of the income year, which is the book-value case analysed in Section 9.2. The distinction between end-of-period and beginning-of-period assessment is thus primarily relevant for *direct* holdings of listed shares.

Our framework nests both cases. End-of-period assessment is the general case analysed throughout the paper, with after-tax discount rate $k^A = (1 - \tau_w)(1 + k) - 1 = k - \tau_w - \tau_w k$. Beginning-of-period assessment is the special case in which the tax base is predetermined (known at the start of the period). When the tax base is deterministic, the cross term $\tau_w k$ vanishes from the pricing equation, and k^A reduces to $k - \tau_w$—the Bjerksund and Schjelderup rate. In both cases, the tax reduces the numerator and denominator of the valuation in the same proportion, and pricing neutrality ($V = V^0$) obtains.

Partial equilibrium. The asset price process $(\boldsymbol{\mu}, \mathbf{V})$ is taken as given. In general equilibrium where all investors face the wealth tax, the return process may in principle adjust. However, under CRRA preferences the optimal portfolio weights are independent of τ_w (Proposition 6b), so that market clearing produces the same equilibrium returns and prices—the partial-equilibrium assumption is self-fulfilling (Section 8.3.3). Departure from partial equilibrium therefore requires either non-homothetic preferences or a violation of the other conditions listed here (universality, market-value tax base).

No borrowing constraints. If the wealth tax reduces the effective return sufficiently that investors wish to lever up, binding leverage constraints create differential effects across wealth levels.

All wealth remains within the tax base. The main text assumes dividends are consumed and the tax is paid by selling shares, which yields the clean proportional-dilution formula $N_n = N_0(1 - \tau_w)^n$. This is a simplifying convention. What matters for neutrality is that all wealth—including reinvested dividends—remains subject to wealth tax at market value. If dividends are reinvested in taxable assets, the multiplicative separability $W_n^A = (1 - \tau_w)^n W_n^B$ is preserved, because the tax factor $(1-\tau_w)$ applies to total wealth regardless of its composition (Appendix A). Neutrality would break if dividends could be channelled into assets that escape the tax base—but that is a violation of universality, not a consequence of reinvestment per se.

Frictionless rebalancing. The multi-asset portfolio results (Propositions 2–6) assume that the investor can rebalance to the optimal weights $\mathbf{w}^*$ at each period end without transaction costs. If rebalancing incurred proportional costs (bid-ask spreads, commissions, or market impact), these costs would reduce wealth in a way that depends on portfolio composition, the degree of weight drift, and the volatility of relative asset returns—introducing a channel through which the wealth tax could interact with portfolio choice indirectly, since a lower wealth base may alter the cost-benefit calculus of rebalancing. In the single-asset case, no rebalancing is needed and this condition is vacuous.

Tax payment mechanism. The main text assumes the tax is paid by selling shares (proportional dilution). If instead the tax is paid from dividend income, the results are preserved when the dividend yield is deterministic, but the multiplicative separability can break down when the dividend yield is stochastic. See Appendix A for a detailed analysis.

Section 9 formalises three channels through which relaxing these conditions produces non-neutral effects: book-value taxation (Section 9.2), liquidity frictions (Section 9.3), and dividend extraction (Section 9.4).

8.2 Independence from Asset Pricing Models

A common source of confusion in discussions of wealth taxation and asset valuation is the role of asset pricing models—in particular, whether results such as those presented here presuppose the Capital Asset Pricing Model (CAPM) or some other equilibrium pricing framework. We address this explicitly, because the distinction between portfolio theory and equilibrium pricing is frequently conflated.

Our results do not assume CAPM. The propositions in this paper are derived from **portfolio theory** (Markowitz, 1952)—the investor's demand-side optimisation problem. We take the return distribution $(\boldsymbol{\mu}, \mathbf{V})$ as exogenous (Assumption A4) and solve for the optimal portfolio weights. This is not an equilibrium statement: we never derive $\boldsymbol{\mu}$ from a market-clearing condition or from a factor model. The expected return vector could come from CAPM, from an APT model, from a multi-factor model, or simply from historical estimation—the results are identical in all cases. What matters is that $(\boldsymbol{\mu}, \mathbf{V})$ exists and does not depend on the tax rate; how it is determined is irrelevant to the analysis.

The distinction is important: CAPM is a statement about equilibrium asset prices (it says the expected excess return on any asset is proportional to its beta with the market portfolio). Our propositions are statements about an individual investor's portfolio problem, conditional on whatever expected returns the market offers. These are logically independent.

Bjerksund and Schjelderup (2022) do not depend on CAPM. Their pricing neutrality result (which corresponds to our Proposition 8) is sometimes misread as requiring CAPM, perhaps because they refer to a "cost of capital" k and to "relevant risk characteristics" of the asset. In fact, their derivation rests on a **no-arbitrage / efficient markets** assumption: in an efficient market, the investor's opportunity return (the discount rate) equals the expected return on an investment with the same risk characteristics as the asset. This is a no-arbitrage condition, not a CAPM-specific one.

CAPM appears in Bjerksund and Schjelderup (2022) only in footnotes—footnote 6 (p. 876) mentions the CAPM formula $r = r_f + \beta(r_M - r_f)$ as one example of how the expected return r could be explained, and footnote 7 (p. 877) notes that "within the capital asset pricing model, for instance, the relevant risk characteristic is measured by beta." Both statements are illustrative, not assumptions. The mechanism behind pricing neutrality—that the tax simultaneously reduces the expected cash flows and the discount rate, and these effects cancel exactly—works for **any** pricing model that provides a discount rate consistent with no-arbitrage.

Summary of what is assumed. Table 3 clarifies the pricing framework used at each level.

Table 3: Pricing frameworks and their role in this paper.

Framework	Provides	Assumes	Used here?
No-arbitrage	Discount rate = opp. return	Efficient markets	Yes (Prop 4′, B&S)
Markowitz	Optimal portfolio weights	$(\boldsymbol{\mu}, \mathbf{V})$ exist; risk aversion	Yes (Props 1′–3′)
CAPM	$\boldsymbol{\mu}$ from equilibrium	Homogeneous expect.	**No**
Multi-factor / APT	$\boldsymbol{\mu}$ from factor loadings	Factor structure	**No**

This distinction is a strength of the analysis: the results hold for any pricing model consistent with no-arbitrage, including CAPM, but they do not require it. The compatibility of our distributional assumptions with CAPM (via the elliptical class) is discussed in Appendix B and in the entry on Hamada and Valdez (2008) in Appendix D.

8.3 The CAPM Special Case

The preceding subsection established that our results do not require CAPM. We now show what happens when CAPM *is* imposed as the equilibrium pricing model. This exercise provides three additional insights: it confirms that the partial-equilibrium results extend to general equilibrium under CRRA preferences, sharpening the caveat in Section 8.1; it identifies a precise error in Fama (2021); and it provides the equilibrium mechanism for the beta-dependent pricing in Section 9.2.4.

8.3.1 After-tax beta invariance

Under the proportional wealth tax, the after-tax return on asset j is $R_j^A = (1-\tau_w)(1+R_j) - 1$ (Equation (33)) and the after-tax market return is $R_M^A = (1-\tau_w)(1+R_M) - 1$. The after-tax beta is:

$$\beta_j^A \;=\; \frac{\mathrm{Cov}(R_j^A,\, R_M^A)}{\mathrm{Var}(R_M^A)} \;=\; \frac{(1-\tau_w)^2\,\mathrm{Cov}(R_j,\, R_M)}{(1-\tau_w)^2\,\mathrm{Var}(R_M)} \;=\; \beta_j. \tag{47}$$

The $(1-\tau_w)^2$ cancels: after-tax beta equals pre-tax beta. This is the CAPM counterpart of the tangency-portfolio invariance in Proposition 6.

8.3.2 The after-tax security market line

Investor B (untaxed) faces the standard SML:

$$E[R_j] \;=\; r_f + \beta_j \cdot \underbrace{\left(E[R_M] - r_f\right)}_{\text{MRP}}. \tag{48}$$

Investor A (taxed) faces the after-tax SML:

$$E[R_j^A] \;=\; r_f^A + \beta_j \cdot \underbrace{\left(E[R_M^A] - r_f^A\right)}_{\text{MRP}^A} \tag{49}$$

where, from (36) and (34),

$$\text{MRP}^A \;=\; (1-\tau_w)\left(E[R_M] - r_f\right) \;=\; (1-\tau_w)\cdot \text{MRP}. \tag{50}$$

The after-tax SML has a lower intercept ($r_f^A < r_f$) and a flatter slope ($(1-\tau_w)\cdot\text{MRP}$), but the same beta. In **gross return** space the relationship simplifies to a uniform scaling:

$$1 + E[R_j^A] \;=\; (1-\tau_w)(1+E[R_j]) \qquad \forall\, j. \tag{51}$$

The after-tax SML is a vertical contraction of the pre-tax SML by the factor $(1-\tau_w)$, mirroring the homothetic contraction of the efficient frontier in Figure 3. Both the capital market line (CML) and the SML contract in the same way: the CML retains the same slope (the Sharpe ratio is preserved by (38)) but shifts to a lower intercept r_f^A; the SML retains the same betas but compresses both intercept and slope by $(1-\tau_w)$. Figure 4 illustrates the geometry.

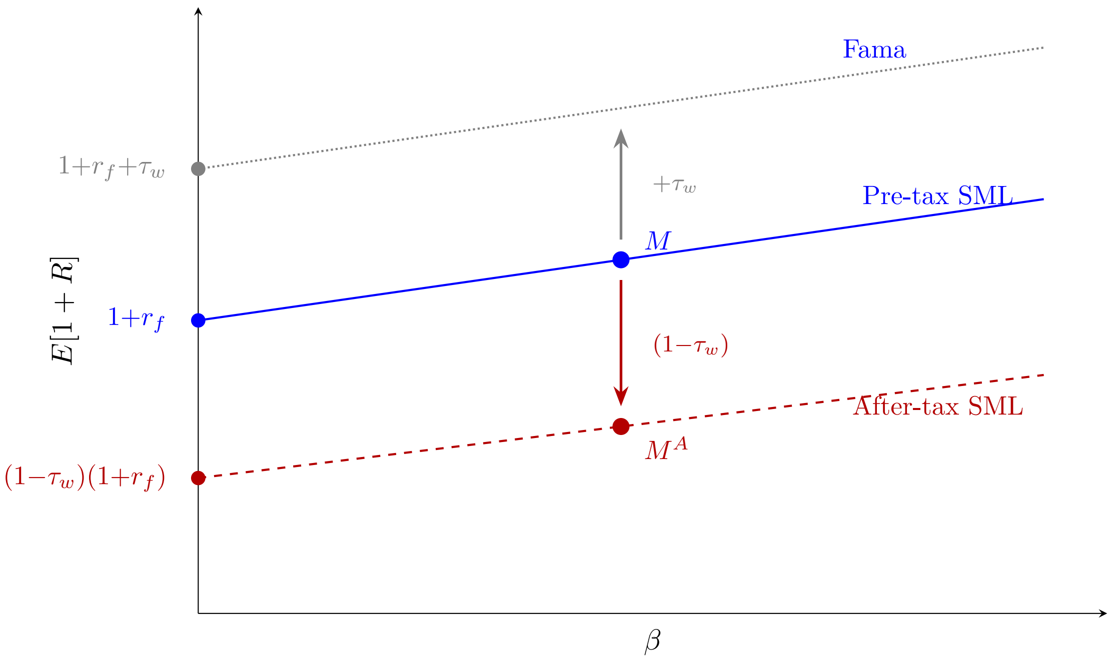

Figure 4: The security market line under a proportional wealth tax. The pre-tax SML (solid blue) connects the risk-free gross return $1{+}r_f$ to the market portfolio M. The correct after-tax SML (dashed red) is a uniform vertical contraction by $(1{-}\tau_w)$: both the intercept and the slope are scaled down, preserving the Sharpe ratio and all betas. Fama's (2021) implicit SML (dotted grey) shifts upward by τ_w, adding the wealth tax to the cost of capital while keeping the risk-free rate unchanged; this overstates the required return because it does not account for the wealth tax on the discount rate itself. Parameters: $r_f = 0.04$, MRP $= 0.06$, $\tau_w = 0.15$ (exaggerated for visibility).

8.3.3 From partial to general equilibrium

The results above are partial equilibrium: the return distribution $(\boldsymbol{\mu}, \mathbf{V})$ is taken as given. We now close the model by imposing CAPM market clearing.

In the standard CAPM derivation, each investor maximises expected utility subject to a budget constraint, and market clearing requires that aggregate demand equals the market capitalisation vector. Under CRRA preferences, Proposition 6(b) establishes that the optimal weight vector $\mathbf{w}^*$ is independent of τ_w. Since the portfolio demand functions are unaffected by the wealth tax, market clearing produces the same equilibrium expected returns and the same asset prices as in the untaxed economy:

$$\boldsymbol{\mu}^{\text{GE, tax}} = \boldsymbol{\mu}^{\text{GE, no tax}}, \qquad \mathbf{V}^{\text{GE, tax}} = \mathbf{V}^{\text{GE, no tax}}. \tag{52}$$

The partial-equilibrium assumption (A4)—that $(\boldsymbol{\mu}, \mathbf{V})$ does not depend on the tax rate—is therefore *self-fulfilling* under CRRA. This sharpens the caveat in Section 8.1: the wealth tax does not disturb general equilibrium when preferences are CRRA and the tax is proportional and universal.

The result extends to a mixed economy with both taxed and untaxed investors. Since both investor types hold the same portfolio weights (Proposition 6), aggregate demand is invariant to the fraction of wealth held by each type. The wealth tax is invisible in CAPM general equilibrium: it transfers a fraction τ_w of end-of-period wealth from investors to the government without distorting any price or allocation.

8.3.4 The error in Fama (2021)

Fama (2021) argues that a wealth tax lowers asset prices. In a Gordon-growth perpetuity, his mechanism is:

$$P_{\text{Fama}} = \frac{D}{r + \tau_w} < \frac{D}{r} = P^0. \tag{53}$$

The wealth tax adds τ_w to the investor's required pre-tax return, raising the effective discount rate from r to $r + \tau_w$. In CAPM terms, this amounts to shifting individual assets upward along a fixed SML—the dotted grey line in Figure 4.

The error is that the SML is not fixed. The wealth tax contracts the *entire* after-tax opportunity set: the risk-free rate falls from r_f to $r_f^A = r_f - \tau_w(1 + r_f)$, and the market risk premium falls from MRP to $(1 - \tau_w) \cdot \text{MRP}$. The investor cannot demand the old after-tax return because every alternative—including the risk-free asset—is also taxed. In the notation of Section 8.4, the correct after-tax discount rate is $k^A = (1 - \tau_w)(1 + k) - 1$, which is *lower* than k, not higher. Both the expected return and the opportunity cost of capital are reduced by the factor $(1 - \tau_w)$, so the ratio—and hence the price—is preserved ($V = V^0$).

Fama's result is internally consistent only if the investor has access to an untaxed outside option—a foreign opportunity set that pins the after-tax required return at the pre-tax level. This is a partial-equilibrium assumption in which the foreign investor's returns are exogenous and the domestic investor is a price-taker against an untaxed benchmark. It produces the same formula as Johnsen and Lensberg (2014): $V = (k - \tau_w)/k \cdot V^0$. When the wealth tax applies to all domestic assets including the risk-free rate, the outside option vanishes, both the numerator and denominator of the valuation are scaled by $(1 - \tau_w)$, and $V = V^0$.

Remark (Scope of the Fama critique). Fama's conclusion that wealth taxes lower prices is not wrong *per se*—it follows from his assumption that asset returns must clear against an untaxed benchmark. The error lies in presenting this as a general equilibrium result when it is, in fact, a partial equilibrium statement: it requires an exogenous opportunity set unaffected by the tax. When the tax is universal, no such benchmark exists, and the price effect vanishes.

Remark (Residence-based wealth taxes and the "untaxed outside option"). A common argument in the policy debate is that, in a small open economy, the marginal price-setting investor is a foreign (untaxed) investor, and that the relevant opportunity set is therefore untaxed. Under this view, a domestic wealth tax raises the required return for domestic investors, who must compete against an untaxed foreign benchmark—exactly Fama's mechanism.

This argument overlooks a fundamental feature of how wealth taxes are implemented in practice. All OECD countries that levy (or have levied) a net wealth tax do so on a *residence basis*: the tax applies to the investor's *worldwide* assets, not merely to domestic holdings (OECD, 2018). A Norwegian investor who buys US equities, German bonds, or any other foreign asset still pays the wealth tax on those holdings. No asset accessible to a domestic resident is exempt—there is no "untaxed outside option" in the investor's opportunity set.

This has three implications. First, the universality condition (Section 8.1) is satisfied automatically for any taxed investor under a residence-based wealth tax on market values: the entire opportunity set—domestic and foreign, risky and risk-free—is taxed at the same rate. Second, the distinction between partial and general equilibrium collapses. The taxed investor's portfolio weights are unchanged (Proposition 6), the untaxed foreign investor's demand is trivially unchanged, and aggregate demand is therefore invariant to the tax—so equilibrium prices do not adjust. For a small open economy, where domestic demand is negligible in global markets, this holds even without the CRRA assumption. Third, even emigration—changing tax residency from taxed to untaxed—does not affect asset prices, because the paper's pricing neutrality (Proposition 8) establishes that both investor types are willing to pay the same price per share. The migration channel affects government revenue but not market prices or portfolio allocations.

Fama's price effect $P = D/(r + \tau_w)$ therefore requires not merely that some investors are untaxed, but that the *taxed* investor can access an asset whose return is not subject to the wealth tax—a condition that no existing or historical OECD wealth tax permits for its residents. The channels through which a wealth tax *does* affect prices are the non-universality channels of Section 9: book-value taxation (which creates heterogeneous effective rates across assets), liquidity frictions, and dividend extraction.

8.3.5 When universality fails: the Security Market Fan

The CAPM framework also illuminates the equilibrium consequences when universality fails. Under book-value taxation (Section 9.2), the effective tax rate on asset j is $\tau_w\theta_j$ where $\theta_j = B_j/V_j$ varies across assets. This creates exactly the heterogeneous-tax setting analysed by Eikseth and Lindset (2009), who show that the SML fans out into a family of lines—one for each effective tax rate—which they term the *Security Market Fan*. In their framework, the after-tax beta becomes

$$\beta_j^\tau = \frac{(1-\tau_j)\,\beta_j}{1-\sum_i w_{i,M}\,\tau_i\,\beta_i} \tag{54}$$

where τ_j is asset j's effective tax rate and $w_{i,M}$ its weight in the market portfolio. For a uniform $\tau_j = \tau_w$ across all assets, the numerator and denominator scale by the same factor and $\beta_j^\tau = \beta_j$, recovering (47). For heterogeneous $\tau_j = \tau_w \theta_j$, assets with low book-to-market ratios (high θ_j) see their after-tax betas compressed, while assets with high book-to-market ratios (low θ_j) retain betas closer to the pre-tax values. The SML fans out, and the equilibrium risk-return relationship is no longer a single line.

This provides the equilibrium mechanism behind the beta-dependent GE pricing formula in Proposition 9(b): book-value taxation creates heterogeneous effective tax rates across assets, which reshapes the SML from a line into a fan and introduces the $(1-\beta)$ term in the valuation formula. Sandvik (2016) makes the same point for Norwegian unlisted shares: the measured beta increases by $1/\theta$ but the risk premium per unit of beta falls by θ, leaving the product for each individual asset unchanged while distorting the cross-sectional ranking.

8.4 The Modigliani-Miller Perspective: The Tax Claim as a Separate Security

The pricing neutrality result (Proposition 8) can be derived from a complementary angle using the Modigliani-Miller framework (Modigliani and Miller, 1958). This perspective treats the government's wealth tax claim as a separate security backed by the firm's cash flows, analogous to how debt is treated in the classical MM analysis of capital structure (Hamada, 1972; Rubinstein, 1973).

Setup. Consider a firm generating a perpetual stream of random cash flows x_t with $E[x_t] = \bar{x}$ and systematic risk β_U. In the absence of taxes, the firm's value and cost of capital are

$$V^0 = \frac{\bar{x}}{k}, \qquad k = r_f + \beta_U \mu,$$

where μ is the market risk premium. Now introduce a proportional wealth tax at rate τ_w on the firm's market value. The government receives $\tau_w V$ per period, where V is the post-tax market value. The firm's pre-tax cash flows are now split between two claimants: the equity holder and the government.

The tax claim has the same risk as the firm. The tax payment $\tau_w V$ is proportional to the firm's market value, so the government's claim has the same systematic risk as the underlying asset: $\beta_{\text{tax}} = \beta_U$. Using the Hamada (1972) leverage formula with the tax claim playing the role of "debt,"

$$\beta_{\text{equity}} = \frac{V^0}{V}\beta_U - \frac{T}{V}\beta_{\text{tax}} = \frac{V^0}{V}\beta_U - \frac{T}{V}\beta_U = \beta_U.$$

The equity beta is unchanged. The government becomes a silent proportional partner, sharing both upside and downside pro rata. Removing a proportional slice does not change the risk profile of what remains.

The discount rate must reflect the tax. The wealth tax contracts the entire after-tax opportunity set by the factor $(1 - \tau_w)$, as established in Section 7.3. Both the risk-free gross return and the risk premium are scaled:

$$1 + r_f^A = (1 - \tau_w)(1 + r_f), \tag{55}$$
$$\mu^A = (1 - \tau_w)\,\mu. \tag{56}$$

The after-tax cost of capital for a β_U-risk asset is therefore

$$k^A = r_f^A + \beta_U\,\mu^A = (1 - \tau_w)(1 + k) - 1. \tag{57}$$

Crucially, $k^A \neq k$: the wealth tax reduces the investor's discount rate together with the expected cash flows.

Pricing neutrality from the MM perspective. The after-tax investor holds a position worth V at the start of the period. At year-end, the total position before tax is $(x + V)$ (dividend plus continuation value), and the tax takes $\tau_w(x + V)$, leaving $(1 - \tau_w)(x + V)$. In a stationary equilibrium the investor reinvests V, so the required-return condition is

$$V(1 + k^A) = (1 - \tau_w)(E[x] + V).$$

Rearranging:

$$V\,(k^A + \tau_w) = (1 - \tau_w)\,E[x] = (1 - \tau_w)\,k\,V^0.$$

Now, $k^A + \tau_w = (1 - \tau_w)(1 + k) - 1 + \tau_w = (1 - \tau_w)\,k$, giving

$$V = \frac{(1 - \tau_w)\,k\,V^0}{(1 - \tau_w)\,k} = V^0. \tag{58}$$

The market price is unchanged. The mechanism is transparent: the factor $(1 - \tau_w)$ appears in both the numerator (reduced after-tax cash flows) and the denominator (reduced opportunity cost of capital), and cancels exactly. This is the valuation counterpart of the orthogonality result (Proposition 7) and confirms Proposition 8 from the MM side.

Comparison with Johnsen and Lensberg (2014). Johnsen and Lensberg model the wealth tax as a separate claim in an MM framework applied to Norwegian listed firms. For the case where the tax base is market value, they derive

$$V_{JL} = V^0\left(1 - \frac{\tau_w}{k}\right). \tag{59}$$

This formula subtracts the tax from the cash flows ($E[x] - \tau_w V^0 = (k - \tau_w)V^0$) but retains the pre-tax discount rate k in the denominator. In their partial equilibrium, where untaxed foreign investors set both the market price and the discount rate, this approach is internally motivated. However, the underlying inconsistency is that the cash flows are reduced by the tax while the discount rate is not:

$$V_{JL} = \frac{\overbrace{(k - \tau_w)}^{\text{after-tax CF}}}{\underbrace{k}_{\text{pre-tax rate}}}\,V^0.$$

Using the after-tax discount rate k^A from (57) for the same after-tax cash flows yields $V = V^0$, as shown in (58). The discrepancy arises because the wealth tax contracts the *entire* after-tax opportunity set—not just the cash flow from the asset in question. Both the expected return and the opportunity cost of capital are reduced by the same factor $(1 - \tau_w)$, so the ratio—and hence the price—is preserved. When the tax base is book value rather than market value, two additional errors in JL's framework emerge; see Section 9.2.5.

Remark (Tax base specification). The neutrality of the MM decomposition depends on the tax being proportional to market value, so that $\beta_{tax} = \beta_U$. If the tax base diverges from market value—for example, book value, assessed value, or lagged historical cost—the tax claim acquires a different beta, and the Hamada (1972) formula produces an equity beta that differs from β_U. In particular, a tax on book value (with $\beta_{\text{tax}} \approx 0$) acts like riskless debt, creating a leverage-like increase in equity beta (cf. Johnsen and Lensberg, 2014, Section 2.3). The tax base specification discussed in Section 8.1 is therefore material for the MM decomposition as well. Section 9.2 refines this observation: $\beta_{\text{tax}} \approx 0$ is correct for the one-period tax obligation, but $\beta_{\text{tax}} = \beta_U$ for the perpetual tax claim under a stationary book-to-market ratio.

8.5 CRRA and the Wealth Effect

For CRRA utility with $\gamma \neq 1$, the value function depends on the wealth level. Since the wealth tax reduces wealth, it can in principle affect portfolio choice through the wealth effect on risk aversion, particularly when the investor optimises consumption jointly with portfolio allocation (the Merton problem; Merton, 1969, 1971). However, the CRRA structure ensures that the optimal *portfolio weights* are invariant to the wealth level, and hence to the tax. This is the content of Proposition 6(b). For non-CRRA preferences (e.g. CARA or habit formation), the wealth effect can break portfolio invariance.

Two supplementary discussions are deferred to appendices: a detailed comparison of the continuous-time and discrete-time formulations (Appendix C) and a survey of the related literature (Appendix D).

9 Beyond Neutrality: Three Channels of Non-Neutral Taxation

9.1 Overview

The neutrality results in Sections 3–7 hold under a specific set of economic conditions, identified in Section 8.1. Each condition, if relaxed, opens a channel through which the proportional wealth tax *does* affect asset prices, portfolio choice, or both. We now formalise three such channels that are of particular practical importance: (i) book-value taxation, (ii) liquidity frictions, and (iii) dividend extraction and investment distortion.

These three channels share a common structural mechanism: each violates the *universality* condition—the requirement that all assets bear the same effective tax rate on market value. When this condition holds, the wealth tax operates as a multiplicative scalar $(1-\tau_w)^n$ that cancels in every ratio, preserving relative returns, Sharpe ratios, and portfolio weights (Section 7.3).

When universality fails, the effective tax rate varies across assets, the multiplicative structure breaks, and the tax distorts relative prices and allocations.

Table 4 summarises the three channels, the condition each relaxes, and the direction of the pricing effect relative to the no-tax value V^0.

Table 4: Three channels of non-neutral wealth taxation.

Channel	Condition relaxed	Direction	Key mechanism
Book value (Section 9.2)	Universality (tax base)	$V > V^0$ when $\theta < 1$	Effective rate $\tau_w \theta_j$ varies
Liquidity (Section 9.3)	Frictionless markets	$V < V^0$	Forced selling at cost c_j
Dividends (Section 9.4)	Frictionless markets	$V < V^0$	Foregone investment to fund tax

The book-value channel is fully formalised using the no-arbitrage pricing framework of Section 8.4. It yields closed-form results for both one-period and multi-period settings, and identifies specific errors in the existing Norwegian literature on wealth tax effects. The liquidity and dividend channels are developed within the same framework, though with more stylised assumptions reflecting the current state of formal knowledge.

9.2 Book-Value Taxation

9.2.1 Motivation and institutional background

In many countries that impose a wealth tax, the tax base deviates from market value for certain asset classes. In Norway, listed shares are taxed at market value, but unlisted shares are taxed at their book value (the firm's equity per share as reported in the tax return), and real estate is taxed at an assessed value that is typically below market value.[1]

Define the *book-to-market ratio*

$$\theta_j \equiv \frac{B_j}{V_j}, \tag{60}$$

where B_j is the book (or assessed) value per share and V_j is the market value. When the wealth tax τ_w is levied on book value rather than market value, the effective tax rate on market value is $\tau_w \theta_j$, which varies across assets. Universality is violated: the tax creates differential costs across assets in exactly the way flagged in Section 8.1.

This section formalises the pricing effect of book-value taxation within the no-arbitrage framework of Section 8.4. As established in Section 8.2, the results follow from no-arbitrage alone; where a general equilibrium extension requires additional structure (CAPM), this is stated explicitly.

The starting point is the Remark on tax base specification in Section 8.4, which observed that a tax on book value (with $\beta_{\text{tax}} \approx 0$) acts like riskless debt in the Hamada (1972) framework. We now show that this observation is correct for the one-period tax obligation but requires refinement for the perpetual tax claim.

[1] The Norwegian wealth tax currently uses approximately 80% of assessed value for primary residences and 100% of assessed value for secondary housing, with assessed values lagging market values. Unlisted shares are valued at their book equity per the company's latest tax balance. Listed shares have been valued at market since 2006.

9.2.2 Three observations

The formalisation rests on three observations about the nature of book-value tax claims.

Observation 1 (Lagged book value is deterministic). Under the Norwegian tax code, the wealth tax liability for year t is based on the book value at the *end of year* $t-1$. At the time the investor makes portfolio decisions for period t, the tax base B_{t-1} is already known. The one-period tax obligation $\tau_w B_{t-1}$ is therefore a deterministic cash flow from the investor's perspective.

Observation 2 (Empirical stationarity of θ). The book-to-market ratio θ is empirically highly persistent. Panel studies of US firms report autocorrelations exceeding 0.95 at horizons up to ten years (Pontiff and Schall, 1998). The residual income model of Ohlson (1995) provides the theoretical underpinning: book value and market value are cointegrated through the clean surplus relation, so θ reverts to a long-run mean. Fractional cointegration tests confirm this for up to 89% of individual firms.[2]

Observation 3 (Gordon steady state as special case). If the firm is in a Gordon steady state with all earnings paid out (no growth), then θ is exactly constant: $\theta = k/\text{ROE}$, where ROE is the return on equity and k is the cost of equity. In this case, the perpetual tax claim $\tau_w B_t = \tau_w \theta V_t$ is proportional to market value with a *constant* ratio θ, and the beta of the tax claim equals the beta of the underlying asset: $\beta_{\text{tax}} = \beta_U$. This contrasts with the one-period case where $\beta_{\text{tax}} \approx 0$.

9.2.3 One-period result

We begin with the one-period case, where the result is model-free.

Consider an asset with random end-of-period payoff x (inclusive of dividends and terminal value), current market value V, and no-tax value V^0. The one-period risk-free rate is r_f. The investor holds the asset for one period and pays wealth tax $\tau_w B$ on the known book value B at period end.[3]

Theorem 1 (One-Period Book-Value Pricing). *Under no-arbitrage, the market value of the asset under book-value wealth taxation is*

$$V = \frac{1}{1-\tau_w}\left(V^0 - \frac{\tau_w B}{1+r_f}\right), \tag{61}$$

provided $\theta < 1 + r_f$.

Proof sketch. The after-tax cash flow is $x - \tau_w B$, where $\tau_w B$ is deterministic by Observation 1. Under the risk-neutral measure $\mathbb{Q}$ (which is unaffected by the introduction of a deterministic

[2] Based on cross-sectional studies using ARFIMA and fractional cointegration models applied to US panel data; see the discussion in Pontiff and Schall (1998) and the references therein.

[3] We write B without a time subscript because it is known at the start of the period; it is the lagged book value from the previous year-end.

liability), the no-arbitrage price satisfies

$$V = \frac{1}{1 + r_f^A} \left(E^{\mathbb{Q}}[x] - \tau_w B \right),$$

where $r_f^A = r_f - \tau_w(1 + r_f)$ is the after-tax risk-free rate from the main text. Since $E^{\mathbb{Q}}[x] = V^0(1 + r_f)$ by definition of the no-tax value, substitution and simplification yield (61). The condition $\theta < 1 + r_f$ ensures $V > 0$; it is satisfied whenever the book value is less than the future value of the market price, which holds empirically for virtually all assets. □

Remark (Generality). Like Proposition 8, this result is distribution-free and holds under any asset pricing model consistent with no-arbitrage. The only additional requirement is that the tax base B is deterministic—which, by Observation 1, is an institutional fact, not a modelling assumption.

Remark (Direction of the effect). When $\theta < 1$ (book value below market value, the empirically dominant case), the taxed asset is worth *more* than under market-value taxation: the investor pays tax on a smaller base, so the effective tax burden is lighter. In the limit $\theta \to 0$ (negligible book value), the tax vanishes entirely and $V \to V^0/(1 - \tau_w)$—the full benefit of a tax-exempt asset. When $\theta = 1$ (book equals market), the formula reduces to $V = V^0$, recovering market-value neutrality.

9.2.4 Multi-period extension

The one-period result extends to a perpetual setting by combining Observations 2 and 3 with the Gordon steady-state model. We consider a firm that pays out all earnings in perpetuity, with cost of equity k and a stationary book-to-market ratio θ. We develop two cases: partial equilibrium (PE), where the discount rate is unaffected by the tax, and general equilibrium (GE), where the market risk premium adjusts because the market portfolio itself is subject to book-value taxation.

Proposition 9 (Multi-Period Book-Value Pricing). *Under no-arbitrage and a stationary book-to-market ratio θ:*

(a) ***Partial equilibrium.*** *If the discount rate k is unaffected by the tax,*

$$V_{\text{PE}} = \frac{k}{k - \tau_w(1 + r_f - \theta)} V^0. \tag{62}$$

(b) ***General equilibrium.*** *If the market risk premium adjusts so that the after-tax required return reflects the systematic risk of the tax claim, and the asset has market beta β,*

$$V_{\text{GE}} = \frac{k}{k - \tau_w(1 - \beta)(1 + r_f - \theta)} V^0. \tag{63}$$

Derivation sketch. The required return under book-value taxation satisfies $k^A = k - \tau_w(1 + r_f) + \tau_w\theta$, where the first two terms are the standard after-tax discount rate from the main text and the last term reflects the *lower* effective tax rate due to the book-value base. The PE pricing follows from $V = E[x]/k^A$ with $E[x] = kV^0$.

For the GE case, the market portfolio is itself subject to book-value taxation. Decomposing the required return into risk-free and risk-premium components, and noting that the risk-free rate adjusts fully (the risk-free asset has $\theta = 1$ if taxed at market value) while the risk premium adjusts only for the $(1 - \beta)$ component that reflects the book-value discount, yields (63). This decomposition uses the CAPM structure to separate the discount rate into its risk-free and systematic risk components—the only point at which an equilibrium pricing model enters. □

Remark (PE vs. GE)*.* The PE formula (62) is beta-independent: it uses only the cost of equity k and the book-to-market ratio θ, and holds under any pricing model. The GE formula (63) introduces beta-dependence because the market risk premium itself adjusts. For $\beta = 1$ (the market portfolio), the GE formula gives $V = V^0$: the book-value benefit is fully absorbed into a lower equilibrium required return, restoring neutrality for the market as a whole. For $\beta < 1$, the asset retains a net benefit ($V > V^0$); for $\beta > 1$, the required return falls by more than the tax benefit, and $V < V^0$. The underlying equilibrium mechanism is the Security Market Fan of Eikseth and Lindset (2009): heterogeneous effective tax rates $\tau_w \theta_j$ across assets reshape the SML from a single line into a fan, producing the beta-dependent pricing; see Section 8.3.5.

Remark (Connection to the tax base Remark)*.* The Remark on tax base specification in Section 8.4 observed that a book-value tax claim with $\beta_{\text{tax}} \approx 0$ acts like riskless debt. This is precisely the one-period result (Theorem 1), where the deterministic tax $\tau_w B$ is discounted at the risk-free rate. The multi-period result refines this: when θ is stationary, the perpetual tax claim has $\beta_{\text{tax}} = \beta_U$ (Observation 3), and the leverage-like effect identified in the Remark disappears for the capitalised claim. The transition from $\beta_{\text{tax}} \approx 0$ (one-period) to $\beta_{\text{tax}} = \beta_U$ (perpetuity) is the key structural difference.

9.2.5 Comparison with Johnsen and Lensberg

Johnsen and Lensberg (2014) analyse the effect of book-value wealth taxation on equity prices within a partial equilibrium CAPM framework. Their central result is a pricing formula

$$V_{JL} = \frac{k - \tau_w}{k} \, V^0,$$

which they use to argue that the wealth tax reduces asset values by approximately 10% (for Norwegian tax rates). This formula, already discussed in Section 8.4, contains two compounding errors.

Corollary 1 (Identification of errors in JL)**.** *The Johnsen and Lensberg (2014) pricing formula contains two errors relative to the no-arbitrage pricing derived in Theorem 1 and Proposition 9:*

(i) ***Pre-tax discount rate.*** *JL discount after-tax cash flows at the pre-tax rate k, whereas no-arbitrage requires the after-tax rate $k^A = k - \tau_w(1 + r_f) + \tau_w\theta$. This double-counts the tax: cash flows are reduced by the tax, but the discount rate does not reflect the corresponding reduction in opportunity cost.*

(ii) ***Risk of the tax claim.*** *JL implicitly treat the tax claim as having $\beta_{\text{tax}} = 0$ (the one-period characterisation), but apply this to a perpetual setting. Under the Gordon steady state with stationary θ, the correct beta is $\beta_{\text{tax}} = \beta_U$ (Observation 3), so the tax claim carries the same systematic risk as the underlying asset. This error is identified by Hansen and Sandvik (2022).*

These errors compound: (i) overstates the tax burden, and (ii) misallocates the risk, producing the spurious prediction that the wealth tax reduces asset values by approximately 10%.

9.2.6 Numerical illustration

Table 5 illustrates the pricing effect for a representative parameterisation: $r_f = 0.03$, $k = 0.10$, $\tau_w = 0.01$, $\theta = 0.50$, and $V^0 = 100$.

Table 5: Book-value wealth tax: pricing effects under alternative models.

Model	V	Comment
No tax	100.0	Benchmark V^0
Market-value tax (neutral)	100.0	Proposition 8
JL formula	90.0	$V_{JL} = (k - \tau_w)/k \cdot V^0$
One-period (Theorem 1)	100.5	$(V^0 - \tau_w B/(1 + r_f))/(1 - \tau_w)$
PE perpetuity (Proposition 9a)	105.6	$kV^0/(k - \tau_w(1 + r_f - \theta))$
GE, $\beta = 0.5$	102.7	(63)
GE, $\beta = 1.0$	100.0	Reduces to V^0
GE, $\beta = 1.5$	97.4	(63)

The sign reversal relative to JL is the most striking feature: JL predict a 10% value *decrease*, while the correct no-arbitrage pricing in partial equilibrium yields a 0.5% to 5.6% value *increase*. The direction reverses because the book-value base ($\theta = 0.50$) creates a lighter effective tax burden than market-value taxation, so assets taxed on book value are worth more, not less, than under the neutral benchmark.

The magnitude depends on the horizon: the one-period effect is small (0.5%) because only the current-year tax obligation benefits from the book-value discount. The perpetuity effect is larger (5.6% in PE) because the entire future stream of tax obligations is discounted, and each one benefits from $\theta < 1$. The GE beta-dependence arises because the market risk premium adjusts: for the market portfolio ($\beta = 1$), the book-value benefit is fully absorbed into a lower required return, so $V = V^0$. For $\beta < 1$ the asset retains a net benefit ($V > V^0$), while for $\beta > 1$ the required return falls by more than the tax benefit, and $V < V^0$.

9.3 Liquidity Frictions

9.3.1 Setup

The main text assumes that the investor can liquidate shares to pay the wealth tax without cost (Section 8.1, "Frictionless rebalancing"). In practice, selling assets to pay a wealth tax incurs transaction costs that vary across asset classes: listed equities face bid-ask spreads and market impact costs; real estate faces agent fees, transfer taxes, and search frictions; private equity positions may be effectively non-tradeable for extended periods.

We model this by introducing a stochastic illiquidity cost c_j for asset j: when the investor sells a fraction of asset j to pay the wealth tax, a fraction c_j of the proceeds is lost to transaction costs. The effective tax rate on market value becomes $\tau_w/(1-c_j)$, because the investor must sell more shares to generate the required after-cost proceeds. Since c_j varies across assets and is stochastic, universality is violated: the effective tax burden differs across assets, and the multiplicative structure that drives neutrality breaks down.

The liquidity channel differs from the book-value channel in an important respect: it operates even when the tax base is market value. Book-value taxation creates non-neutrality through the *definition* of the tax base; liquidity frictions create non-neutrality through the *mechanics* of tax payment. Both channels violate universality, but through different mechanisms.

Remark (Practical payment mechanics). The proportional-dilution mechanism is a modelling abstraction that represents an upper bound on the liquidity effect. In practice, the wealth tax is rarely paid by selling shares in a single year-end transaction. In Norway, the personal wealth tax (assessed on 31 December values) is collected through quarterly advance payments (*forskuddsskatt*) during the assessment year, smoothing the liquidity demand over time. More importantly, investors typically cover the tax through a hierarchy of payment sources: (i) other liquid assets or income (salary, interest, rental income), (ii) dividend payments from the company, and (iii) share sales—with outright forced liquidation as a last resort rather than the default mechanism. Empirical evidence supports this hierarchy: Berzins et al. (2022) find that Norwegian private firms increase dividend payouts in response to their owners' wealth tax obligations, while Ring (2024) documents precautionary saving behaviour consistent with households maintaining liquid buffers to meet future tax liabilities. Alstadsæter et al. (2022) observe that countries implementing wealth taxes make "practical compromises regarding treatment and valuation of different asset categories to ease assessment and liquidity difficulties." The forced-selling model should therefore be interpreted as characterising the marginal cost of the *least liquid* assets in the portfolio—those for which the payment hierarchy is exhausted—rather than the average cost across all holdings. The dividend extraction channel (Section 9.4) formalises the more common payment mechanism.

9.3.2 Multi-factor neutrality without friction

Before introducing frictions, we note that the neutrality results of Sections 3–7 extend immediately to any multi-factor pricing model. Whether expected returns are determined by a single market factor (CAPM), three factors (Fama–French), or an arbitrary K-factor model with factors $f_1, \dots, f_K$ (which may include a liquidity factor), the wealth tax operates as a multiplicative scalar that reduces all factor premia uniformly:

$$E[R_j^A] - r_f^A = (1-\tau_w)\big(E[R_j^B] - r_f\big) = (1-\tau_w)\sum_{k=1}^{K} \beta_{j,k}\,\lambda_k,$$

where λ_k is the premium for factor k. The factor betas $\beta_{j,k}$ are unchanged, and the relative premia across assets are preserved. Adding liquidity as a priced factor—as in the models of Acharya and Pedersen (2005), Pástor and Stambaugh (2003), or Amihud (2002)—changes nothing about neutrality, as long as the tax is paid without friction. The key insight is that neutrality is a property of the *tax mechanism* (multiplicative scaling), not of the asset pricing model.

9.3.3 Pricing with friction

Introduce a stochastic illiquidity cost $c_j \in [0, 1)$ that the investor incurs when liquidating shares of asset j to pay the wealth tax. The cost c_j may be correlated with the asset's payoff x_j and with aggregate market conditions.

The after-tax cash flow from one share of asset j is

$$x_j - \tau_w V \cdot \frac{1}{1 - c_j},$$

where the factor $1/(1-c_j)$ reflects the additional shares that must be sold to cover the transaction cost. Since c_j is stochastic, the after-tax factor $(1 - \tau_w/(1 - c_j))$ is random, and the clean multiplicative structure of the main text is lost.

Proposition 10 (Pricing with Liquidity Frictions)**.** *Under no-arbitrage, the one-period price of asset j under wealth taxation with stochastic illiquidity cost c_j satisfies*

$$V \approx V^0 - \frac{\tau_w}{1 + r_f} V \cdot E^{\mathbb{Q}}[c_j] - \frac{\tau_w}{(1 + r_f)^2} \mathrm{Cov}^{\mathbb{Q}}(x_j + V,\ c_j), \tag{64}$$

where $E^{\mathbb{Q}}$ and $\mathrm{Cov}^{\mathbb{Q}}$ denote expectation and covariance under the risk-neutral measure.

Derivation sketch. The no-arbitrage price satisfies $V = E^{\mathbb{Q}}[x_j - \tau_w V/(1-c_j)]/(1+r_f^A)$. Expanding $1/(1 - c_j) \approx 1 + c_j$ for small c_j and separating the deterministic and stochastic components yields three terms: (i) the no-tax value V^0 (from the main text neutrality result), (ii) a level effect proportional to $E^{\mathbb{Q}}[c_j]$ representing the expected illiquidity cost, and (iii) a covariance effect capturing the risk premium from the correlation between asset payoffs and illiquidity costs. □

Remark (Two effects)*.* The pricing impact decomposes into a *level effect* and a *covariance effect*:

- **Level effect** ($E^{\mathbb{Q}}[c_j]$): The expected illiquidity cost reduces the asset's value in proportion to the tax rate. This is a first-order effect that depends on the asset's average liquidity.
- **Covariance effect** ($\mathrm{Cov}^{\mathbb{Q}}(x_j + V, c_j)$): If the asset becomes more illiquid precisely when its payoff is low (positive covariance between c_j and losses), the tax-induced forced selling is most costly in bad states of the world. This creates a liquidity risk premium that is specific to the wealth tax mechanism.

Remark (Structural insight)*.* The wealth tax converts a multiplicative (neutral) tax mechanism into a multiplicative-plus-additive mechanism. In the frictionless case, the after-tax factor $(1 - \tau_w)$ is a constant multiplier that cancels in ratios. With friction, the effective factor $(1 - \tau_w/(1 - c_j))$ is stochastic, and its covariance with asset payoffs introduces an additive, asset-specific distortion that does not cancel.

9.3.4 Connection to the liquidity pricing literature

The covariance effect in (64) maps directly onto the three liquidity betas identified in the liquidity-adjusted CAPM of Acharya and Pedersen (2005). In their framework, expected returns reflect not only the covariance of returns with the market, but also three additional terms:

$$\beta_{c,c}: \quad \text{Cov}(c_j, c_M), \tag{65}$$

$$\beta_{c,r}: \quad \text{Cov}(c_j, R_M), \tag{66}$$

$$\beta_{r,c}: \quad \text{Cov}(R_j, c_M), \tag{67}$$

where c_M denotes aggregate market illiquidity and R_M is the market return. The first captures commonality in liquidity, the second captures the tendency for assets to become illiquid when the market rises (or falls), and the third captures the tendency for returns to be low when the market is illiquid.

The wealth tax *amplifies* all three channels. Forced selling to pay the tax creates correlated liquidity demand across all taxed investors, increasing commonality in illiquidity costs ($\beta_{c,c}$). If tax-induced selling is concentrated in downturns (when portfolio values are lower relative to tax liabilities), it increases the covariance between illiquidity and market conditions ($\beta_{c,r}$ and $\beta_{r,c}$). The magnitude depends on the fraction of aggregate ownership subject to the wealth tax and on the concentration of tax-motivated selling in time.

Remark. The liquidity channel operates even under market-value taxation. Unlike the book-value channel, which requires a divergence between the tax base and market value, the liquidity channel arises from the mechanics of tax *payment*: any tax that must be paid by selling assets incurs transaction costs. The effect is therefore a general feature of wealth taxation, not specific to any particular tax base specification.

Pástor and Stambaugh (2003) estimate a tradeable liquidity factor with a risk premium of approximately 7.5% per year. Amihud (2002) documents a significant cross-sectional relationship between illiquidity (measured as price impact per unit of trading volume) and expected returns. The wealth tax mechanism described here provides a specific channel through which these liquidity premia are amplified: the tax creates a periodic, predictable source of forced selling that is correlated across investors and concentrated in time.

9.4 Dividend Extraction and Investment Distortion

9.4.1 Relation to Appendix A

Appendix A analyses the financial mechanics of dividend payment vs. share sales and shows that the *method* of tax payment is irrelevant for pricing neutrality (Appendix A.5). The present section addresses a different question: when the wealth tax forces the firm (or investor) to extract dividends that would otherwise have been reinvested, the resulting loss of investment opportunities has a real cost that breaks pricing neutrality. The non-neutrality arises not from the financial mechanics of dividend payment, but from the *foregone investment* that the dividend extraction necessitates.

9.4.2 Tax payment constraint

Consider a firm with cost of equity k, payout ratio δ, and expected cash flow $\bar{x} = kV^0$. Under book-value taxation, the wealth tax liability per share is $\tau_w \theta V$ per period. If the investor cannot (or prefers not to) sell shares to pay the tax, the dividend must cover the tax liability, imposing a minimum payout constraint:

$$\delta \geq \frac{\tau_w \theta}{k}. \tag{68}$$

This constraint binds when the firm would otherwise choose a lower payout ratio—typically growth firms with high reinvestment needs—and when the investor faces high costs of selling shares, as is the case for private firms with illiquid equity.

For a concrete illustration: with $\tau_w = 0.01$, $\theta = 0.50$, and $k = 0.10$, the minimum payout ratio is $\delta \geq 0.05$, or 5% of expected cash flows. For a growth firm that would optimally pay out only 2–3% (retaining the rest for reinvestment), the wealth tax forces an increase in the payout ratio that reduces available investment capital.

9.4.3 Investment distortion

When the payout constraint (68) binds, the firm foregoes investment opportunities whose internal rate of return exceeds the cost of capital. The value loss depends on the profitability of the foregone investment.

Proposition 11 (Investment Distortion)**.** *If the firm has investment opportunities with internal rate of return $\rho > k$, and the payout constraint* (68) *binds, the value loss relative to no tax is*

$$\Delta V = \frac{\rho - k}{k} \left(\tau_w \theta \, V - \delta \, \bar{x} \right), \tag{69}$$

where $\delta \, \bar{x}$ is the dividend the firm would pay absent the tax constraint, and $\tau_w \theta \, V - \delta \, \bar{x}$ is the additional payout forced by the tax.

Derivation sketch. The forced additional payout $\tau_w \theta \, V - \delta \, \bar{x}$ is redirected from investment (earning ρ) to the tax authority. Each unit of foregone investment destroys present value of $(\rho - k)/k$ in perpetuity. The total value loss is the product of the unit cost and the volume of displaced investment. □

Remark. The distortion is larger when (i) the firm has high growth opportunities ($\rho \gg k$), (ii) the book-to-market ratio θ is high (increasing the tax burden on book value), and (iii) the firm's optimal payout ratio δ is low (so the constraint binds more tightly). The distortion vanishes for firms whose optimal payout already exceeds the tax threshold—typically mature firms with limited growth opportunities.

9.4.4 Optimal payment mechanism

The investor faces a trade-off between two mechanisms for paying the wealth tax: selling shares at illiquidity cost c_j, or extracting dividends at investment distortion cost $(\rho - k)/k$. The optimal mechanism minimises the total cost:

$$\text{Cost}_{\text{sell}} = c_j \cdot \tau_w \theta V, \qquad \text{Cost}_{\text{dividend}} = \frac{\rho - k}{k} \cdot \tau_w \theta V. \tag{70}$$

Selling is preferred when $c_j < (\rho - k)/k$ (low illiquidity cost relative to growth opportunities), and dividend extraction is preferred when $c_j > (\rho - k)/k$ (high illiquidity cost, limited growth).

This creates a natural partition across asset classes:

- **Listed, mature firms:** Low c_j, low ρ. Sell shares to pay tax; minimal distortion.
- **Listed, growth firms:** Low c_j, high ρ. Sell shares; moderate distortion from market impact.
- **Private, mature firms:** High c_j, low ρ. Extract dividends; moderate distortion.
- **Private, growth firms:** High c_j, high ρ. Both mechanisms are costly; maximal distortion. This is the asset class most affected by the wealth tax.

Berzins et al. (2022) provide direct evidence for this partition using Norwegian firm-level data. They find that the wealth tax leads to significantly higher dividend payouts in private firms, with the effect concentrated among firms with controlling shareholders facing binding tax obligations. The associated reduction in investment is economically large: firms with higher wealth tax exposure invest less and grow more slowly, consistent with the dividend extraction channel formalised here.

9.5 Combined Effects and Implications

9.5.1 Interaction of the three channels

The three channels identified above have opposing effects on asset values. Book-value taxation raises values ($V > V^0$) when the tax base is below market value ($\theta < 1$), because the effective tax burden is lighter. Liquidity frictions and dividend extraction both lower values ($V < V^0$), because they impose real costs on the tax payment process.

The net effect is asset-class dependent. Table 6 summarises the qualitative direction for the principal asset classes affected by the Norwegian wealth tax.

Table 6: Net pricing effect across asset classes (qualitative).

Asset class	θ	c_j	Book value	Liquidity	Dividend
Listed, mature	0.5–1.0	Low	+small	−small	−negligible
Listed, growth	0.2–0.5	Low	+moderate	−small	−small
Private, mature	0.5–1.0	High	+small	−moderate	−moderate
Private, growth	0.2–0.5	High	+moderate	−large	−large
Real estate	0.3–0.8	High	+moderate	−moderate	N/A

For listed assets with low illiquidity costs, the book-value discount ($\theta < 1$) dominates, and the net effect is likely a modest value *increase*. For private firms with high illiquidity costs and significant growth opportunities, the liquidity and dividend channels dominate, and the net effect is a value *decrease*—potentially substantial.

9.5.2 Empirical implications

The framework developed in this section explains several empirical observations that are difficult to reconcile with either pure neutrality or the predictions of Johnsen and Lensberg (2014):

Real effects concentrated in private firms. Berzins et al. (2022) document that the wealth tax leads to higher dividends and lower investment in private firms, with negligible effects on listed firms. This is consistent with the asset-class partition above: listed firms have low illiquidity costs and can sell shares to pay the tax, while private firms face high c_j and must extract dividends.

Book-value discount as partial offset. The Norwegian practice of valuing unlisted shares at book value—often 50% or less of market value—mitigates the liquidity and dividend costs by reducing the effective tax rate. The book-value channel provides a *partial offset* to the real costs imposed by the other two channels. This explains why the observed effects, while significant, are smaller than a naïve calculation based on the statutory wealth tax rate would suggest.

Precautionary saving. Ring (2024) finds evidence that the wealth tax induces precautionary saving, with households maintaining higher liquid asset balances to ensure they can meet future tax obligations. This is consistent with the liquidity channel: the expected cost of forced selling creates an incentive to hold more liquid assets, even at the expense of lower expected returns.

9.5.3 Policy implications

The analysis yields several insights for tax policy design. First, the welfare cost of a wealth tax is not uniform across asset classes. The distortionary effects are concentrated among illiquid assets and growth firms, while liquid, mature assets are largely unaffected.

Second, a market-value tax base is more neutral than a book-value base, in the specific sense that it eliminates the asset-specific variation in effective tax rates that the book-value channel creates. However, this does not mean that market-value taxation eliminates all non-neutrality: the liquidity and dividend channels operate regardless of the tax base specification.

Third, the Norwegian practice of valuation discounts for unlisted assets and real estate has the effect of partially compensating for the real costs imposed by illiquidity and dividend extraction. While this compensation is not formally calibrated to the liquidity costs, it operates in the right direction: assets with the highest illiquidity costs (private firms, real estate) receive the largest valuation discounts ($\theta \ll 1$).

Fourth, the interaction between the channels suggests that a well-designed wealth tax would explicitly account for illiquidity costs in setting valuation discounts, rather than relying on historical book values that may or may not approximate the relevant cost. This would require estimating asset-specific illiquidity costs c_j—a substantial empirical undertaking, but one that would improve the alignment between the effective tax burden and the actual cost of compliance.

10 Conclusion and Further Work

This paper establishes that a proportional wealth tax levied at a uniform rate on the market value of all assets is neutral with respect to four dimensions of the investor's problem: the risk-reward profile of wealth (CV invariance), the optimal portfolio weights (including the tangency portfolio), the geometric relationship between the tax and portfolio choice (orthogonality), and the per-share asset price (pricing neutrality). These results are derived first under geometric Brownian motion and then generalised to any return distribution in the location-scale family; the two distributional assumptions that matter—finite second moments and proportionality of the tax—are minimal.

The neutrality rests on two conditions: universal taxation at market value, and frictionless markets. The payment mechanism—whether the tax is funded from dividends, share sales, or other taxable wealth—does not affect the result, provided that reinvested dividends remain within the tax base (Appendix A). A complementary analysis under the CAPM (Section 8.3) confirms that after-tax betas equal pre-tax betas, the security market line contracts uniformly, and—under CRRA preferences—general equilibrium returns and prices are identical to the no-tax benchmark. This analysis also identifies a pricing error in Fama (2021), who adds the wealth tax to the cost of capital without recognising that the entire opportunity set, including the risk-free rate, is contracted by the same factor.

When these conditions fail, three channels of non-neutrality emerge with opposing price effects: book-value taxation raises valuations for assets with $\theta < 1$; liquidity frictions and dividend extraction lower them. The net effect is asset-class specific and depends on institutional details—book-to-market ratios, illiquidity costs, and the profitability of foregone investment.

Directions for further work. Several extensions would strengthen and complement the analysis.

On the *theoretical* side, the model could be embedded in a multi-period consumption-saving framework with endogenous labour supply, where the wealth tax interacts with the intertemporal allocation of resources. A general equilibrium analysis with heterogeneous agents facing different effective tax rates—as arises under progressive taxation or asset-class exemptions—would connect the Security Market Fan mechanism of Section 8.3.5 to equilibrium asset pricing with realistic tax codes. The interaction between the wealth tax and income or capital gains taxes is unexplored in our framework and relevant for optimal tax design.

On the *empirical* side, four avenues stand out. First, calibrating asset-class-specific illiquidity costs c_j (listed equities, private firms, real estate) would allow a quantitative assessment of the liquidity channel. Second, Norwegian registry data, which link individual portfolio holdings to tax records, could be used to test whether portfolio weights are indeed invariant to the wealth tax rate—exploiting cross-sectional variation in effective rates created by valuation discounts. Third, the book-value channel can be estimated using variation in θ across Norwegian firms combined with the pricing formulas of Proposition 9. Fourth, the dividend-extraction channel predicts that firms whose owners face binding payout constraints exhibit higher payout ratios around tax dates—a prediction testable with payout data.

A note on structure. To streamline the main argument, three self-contained discussions have been placed in appendices: the detailed classification of distributional assumptions (Appendix B), the continuous-time versus discrete-time comparison (Appendix C), and the survey of related literature (Appendix D). These can be read independently without disrupting the main narrative.

Acknowledgements

The author acknowledges the use of Claude (Anthropic) for assistance with literature review, LaTeX typesetting, mathematical exposition, and editorial refinement, and Lemma (Axiomatic AI) for review and proof checking. All substantive arguments, economic reasoning, and conclusions are the author's own.

References

Viral V. Acharya and Lasse Heje Pedersen. Asset pricing with liquidity risk. *Journal of Financial Economics*, 77(2):375–410, 2005.

Annette Alstadsæter, Marie Bjørneby, Wojciech Kopczuk, Simen Markussen, and Knut Røed. Saving effects of a real-life imperfectly implemented wealth tax. *AEA Papers and Proceedings*, 112:131–135, 2022.

Yakov Amihud. Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, 5(1):31–56, 2002.

Janis Berzins, Øyvind Bøhren, and Bogdan Stacescu. Shareholder illiquidity and firm behavior: Financial and real effects of the personal wealth tax. *Review of Finance*, 26(5):1165–1203, 2022.

Petter Bjerksund and Guttorm Schjelderup. Investor asset valuation under a wealth tax and a capital income tax. *International Tax and Public Finance*, 29:873–889, 2022.

Marie Bjørneby, Simen Markussen, and Knut Røed. An imperfect wealth tax and employment in closely held firms. *Economica*, 90(358):557–583, 2023. doi: 10.1111/ecca.12456.

Gary Chamberlain. A characterization of the distributions that imply mean-variance utility functions. *Journal of Economic Theory*, 29:185–201, 1983.

George M. Constantinides. Capital market equilibrium with personal tax. *Econometrica*, 51(3): 611–636, 1983.

Robert M. Dammon, Chester S. Spatt, and Harold H. Zhang. Optimal consumption and investment with capital gains taxes. *Review of Financial Studies*, 14(3):583–616, 2001.

Hans Marius Eikseth and Snorre Lindset. A note on capital asset pricing and heterogeneous taxes. *Journal of Banking & Finance*, 33(3):573–577, 2009.

Eugene F. Fama. Wealth taxes. Paper 20-34, Chicago Booth, 2021.

Mahmoud Hamada and Emiliano A. Valdez. CAPM and option pricing with elliptically contoured distributions. *Journal of Risk and Insurance*, 75(2):387–409, 2008.

Robert S. Hamada. The effect of the firm's capital structure on the systematic risk of common stocks. *Journal of Finance*, 27(2):435–452, 1972.

Pål E. Hansen and Bjørn Sandvik. Formuesskatt på aksjer—replikk til Johnsen og Lensberg. *Samfunnsøkonomen*, 136(1):54–66, 2022.

Jonathan E. Ingersoll. *Theory of Financial Decision Making*. Rowman & Littlefield, 1987.

Thore Johnsen and Terje Lensberg. A note on the cost of collecting wealth taxes. Working Paper, Norwegian School of Economics, 2014.

Lutz Kruschwitz, Andreas Löffler, Daniela Lorenz, and Moritz Uttscheid. Costs of capital and wealth tax: Remarks on Bjerksund and Schjelderup (2022). Discussion Paper 279, arqus, 2023.

Harry Markowitz. Portfolio selection. *Journal of Finance*, 7(1):77–91, 1952.

Robert C. Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *Review of Economics and Statistics*, 51(3):247–257, 1969. doi: 10.2307/1926560.

Robert C. Merton. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413, 1971. doi: 10.1016/0022-0531(71)90038-X.

Jack Meyer. Two-moment decision models and expected utility maximization. *American Economic Review*, 77(3):421–430, 1987.

Franco Modigliani and Merton H. Miller. The cost of capital, corporation finance and the theory of investment. *American Economic Review*, 48(3):261–297, 1958.

OECD. *The Role and Design of Net Wealth Taxes in the OECD*. Number 26 in OECD Tax Policy Studies. OECD Publishing, Paris, 2018. doi: 10.1787/9789264290303-en.

James A. Ohlson. Earnings, book values, and dividends in equity valuation. *Contemporary Accounting Research*, 11(2):661–687, 1995.

Joel Owen and Ramon Rabinovitch. On the class of elliptical distributions and their applications to the theory of portfolio choice. *Journal of Finance*, 38(3):745–752, 1983.

Ľubos Pástor and Robert F. Stambaugh. Liquidity risk and expected stock returns. *Journal of Political Economy*, 111(3):642–685, 2003.

Jeffrey Pontiff and Lawrence D. Schall. Book-to-market ratios as predictors of market returns. *Journal of Financial Economics*, 49(2):141–160, 1998.

Marius A. K. Ring. Wealth taxation and household saving: Evidence from assessment discontinuities in Norway. *Review of Economics and Statistics*, 2024. doi: 10.1093/restud/rdae091. Forthcoming; SSRN Working Paper, 2020.

Mark E. Rubinstein. A mean-variance synthesis of corporate financial theory. *Journal of Finance*, 28(1):167–181, 1973.

Agnar Sandmo. The effects of taxation on savings and risk taking. In *Handbook of Public Economics*, volume 1, pages 265–311. 1985.

Bjørn Sandvik. Formuesskatt på aksjer. *Samfunnsøkonomen*, 130(3):36–47, 2016.

William F. Sharpe. Mutual fund performance. *Journal of Business*, 39(1):119–138, 1966.

Joseph E. Stiglitz. The effects of income, wealth, and capital gains taxation on risk-taking. *Quarterly Journal of Economics*, 83(2):263–283, 1969.

John D. Stowe. Wealth taxes and capital markets. SSRN Working Paper No. 3542463, 2021.

A Alternative Tax Payment Mechanisms: Dividends vs. Share Sales

The main results (Sections 3–7) assume that the wealth tax is paid by selling a fraction τ_w of the investor's position at each period end—the proportional-dilution mechanism. This appendix considers an alternative: the wealth tax is paid from dividend income rather than from share sales. We retain all other assumptions of the main text, including the absence of income and capital gains taxes.

A.1 Setup

At each period end i, the asset pays a dividend D_i per share. The dividend yield is $\delta_i \equiv D_i/P_i$. The wealth tax liability is $\tau_w N_{i-1} P_i$, and the investor uses dividend income $D_i N_{i-1}$ to pay as much of this liability as possible before resorting to share sales.

Three cases arise, depending on the relationship between the dividend yield and the tax rate.

A.2 Case 1: Dividends Always Sufficient ($\delta_i \geq \tau_w$ with Certainty)

If the dividend yield exceeds the tax rate in every period, the tax is fully paid from dividends and no shares are sold:

$$N_n = N_0 \quad \text{(constant)}, \qquad W_n^A = N_0 P_n = W_n^B. \tag{71}$$

The equity wealth of the taxed investor is **identical** to that of the untaxed investor. All four propositions hold trivially—not because of multiplicative separability, but because the tax does not touch the capital base at all.

The economic content is different from the proportional-dilution case. Under proportional dilution, the tax erodes the number of shares but the investor consumes the full dividend D_i per share. Under dividend payment, the investor retains all shares but consumes only $(D_i - \tau_w P_i)$ per share. The tax has effectively become a **consumption tax on dividend income**: it reduces the investor's standard of living without affecting the trajectory of equity wealth.

This distinction matters for welfare analysis (the investor's lifetime utility from consumption is lower) but not for the portfolio choice, risk-reward, or pricing results, which concern equity wealth and asset returns.

A.3 Case 2: Dividends Insufficient, Constant Yield ($\delta < \tau_w$)

If the dividend yield is constant at δ (dropping the time subscript) and below the tax rate, all dividends go to tax and the shortfall $(\tau_w - \delta)P_i$ per share is covered by selling shares. The share count evolves as:

$$N_i = N_{i-1}\big(1 - (\tau_w - \delta)\big), \qquad N_n = N_0\big(1 - (\tau_w - \delta)\big)^n \tag{72}$$

and equity wealth is:

$$W_n^A = N_0\big(1 - (\tau_w - \delta)\big)^n P_n = \big(1 - (\tau_w - \delta)\big)^n \cdot W_n^B. \tag{73}$$

This is exactly the proportional-dilution model of Sections 3–7 with the **effective tax rate** $\tau_w^{\text{eff}} = \tau_w - \delta$ replacing τ_w. The multiplicative separability is preserved, and all four propositions hold with τ_w^{eff} in place of τ_w. The dividend acts as a partial shield against capital erosion: the higher the yield, the smaller the effective dilution rate.

In the limiting case $\delta = \tau_w$, the effective rate is zero and we recover Case 1. In the limiting case $\delta = 0$ (no dividends), we recover the original proportional-dilution model.

A.4 Case 3: Stochastic Dividend Yield

If the dividend yield $\delta_i = D_i/P_i$ varies stochastically across periods, the effective liquidation rate at each period is $\max(\tau_w - \delta_i,\, 0)$, and the share count evolves as:

$$N_n = N_0 \prod_{i=1}^{n} \big(1 - \max(\tau_w - \delta_i,\, 0)\big). \tag{74}$$

The product $\prod_i(1 - \max(\tau_w - \delta_i, 0))$ is now a **stochastic** quantity—it depends on the realised path of dividend yields, which are themselves functions of the price process. The deterministic multiplicative scalar $(1-\tau_w)^n$ of the proportional-dilution model is replaced by a random variable that is correlated with the cumulative return $G^{(n)} = P_n/P_0$.

This breaks the multiplicative separability $W_n^A = c \cdot W_n^B$ with deterministic c, because the tax factor and the return are no longer independent. Specifically:

$$W_n^A = N_0 P_n \prod_{i=1}^{n} \big(1 - \max(\tau_w - \delta_i,\, 0)\big). \tag{75}$$

The product $P_n \cdot \prod_i(1 - \max(\tau_w - \delta_i, 0))$ is a product of correlated random variables, and in general:

$$E[W_n^A] \neq \left(\prod_{i=1}^{n} E[1 - \max(\tau_w - \delta_i, 0)]\right) \cdot E[W_n^B], \tag{76}$$

$$\mathrm{SD}(W_n^A) \neq \left(\prod_{i=1}^{n} E[1 - \max(\tau_w - \delta_i, 0)]\right) \cdot \mathrm{SD}(W_n^B). \tag{77}$$

The coefficient of variation of W_n^A is therefore **not** in general equal to that of W_n^B, and Proposition 5 (CV invariance) may fail. Similarly, the portfolio choice results (Propositions 6 and 7) may be affected, because the stochastic tax factor introduces an additional source of portfolio-level risk that depends on the interaction between dividend yields and asset returns.

A.5 Pricing Neutrality Under All Cases

Proposition 8 (pricing neutrality) is robust to the payment mechanism. The value of a share to the investor depends on the discounted stream of after-tax cash flows. Regardless of whether the tax is paid from dividends or from share sales, the total tax liability is $\tau_w \times$ market value—the method of payment is a financing decision, not a valuation one. In partial equilibrium, the taxed investor's discount rate adjusts to reflect the tax, as shown by Bjerksund and Schjelderup (2022), and the NPV of the share remains equal to the market price for both investors.

More concretely: if the investor pays the tax from dividends, the after-tax dividend per share is $(D_i - \tau_w P_i)$ instead of D_i, but the discount rate also falls from r_f to $r_f^A = r_f - \tau_w(1 + r_f)$. These effects cancel in the discounted cash flow calculation, preserving pricing neutrality. Note, however, that this analysis concerns the *financial mechanics* of tax payment. When the tax forces dividend extraction that displaces profitable investment, a *real* cost arises that breaks pricing neutrality; see Section 9.4.

A.6 Summary

Table 7: Effect of tax payment mechanism on results.

Payment mechanism	N_n	Mult. sep.	Props 1′–3′	Prop 4′
Proportional dilution (share sales)	$N_0(1-\tau_w)^n$, det.	Yes	Hold	Holds
Dividends suff. ($\delta \geq \tau_w$)	N_0, const.	Trivially	Trivially	Holds
Div. insuff., const. $\delta < \tau_w$	$N_0(1-\tau_w^{\text{eff}})^n$, det.	Yes	Hold (adj.)	Holds
Stochastic div. yield	Path-dep., stoch.	**Breaks**	**May fail**	Holds

The proportional-dilution assumption of the main text is not merely a modelling convenience. It ensures that the tax factor $(1 - \tau_w)^n$ is deterministic and independent of the return realisation, which is the structural property that drives Propositions 1′–3′. Alternative payment mechanisms preserve the results if and only if the effective dilution rate remains deterministic. Among the cases considered, this holds when the dividend yield is either zero (original model), constant, or always above the tax rate. When the dividend yield is stochastic and sometimes insufficient, the interaction between the tax payment and the return process introduces a new source of risk that breaks the multiplicative structure.

B The Role of Distributional Assumptions

The location-scale family serves a specific and limited role in the generalisation. It is worth being precise about what it contributes and where it is not needed.

What the location-scale family provides. The assumption contributes two things. First, *two-moment sufficiency*: Meyer (1987) establishes that if all alternatives belong to the same location-scale family, then expected utility can be expressed as $E[U(X)] = V(\mu_X, \sigma_X)$, justifying the mean-variance framework without assuming quadratic utility; see also Ingersoll (1987, Ch. 4). Second, *closure under linear combination*: in the multivariate case (elliptical distributions), any portfolio return inherits the distributional family of the individual asset returns (Hamada and Valdez, 2008, Theorem 2), so that $R_P = \mathbf{w}^\top \mathbf{R}$ has the same shape parameter as $\mathbf{R}$.

These two services are needed for Propositions 2′ and 3′, which require the (σ, μ) representation of the efficient frontier and the geometric contraction argument. They are *not* needed for CV invariance (Proposition 1′), which follows from scalar multiplication alone; for tangency portfolio invariance (Proposition 2′a), which is an algebraic identity; or for pricing neutrality (Proposition 4′), which is entirely distribution-free.

The elliptical hierarchy. The relevant distributional classes form a hierarchy:

Table 8: Distributional classes and their properties.

Class	Univariate examples	Multivariate closure	Portfolio separation
Normal / GBM	Normal, lognormal	Yes	Yes
Elliptical	Student-t, logistic, Laplace	Yes	Yes
Location-scale	All symmetric unimodal	Requires elliptical	Requires elliptical
Finite 2nd moments	Any with $E[R^2] < \infty$	Not necessarily	SR invariance only

Portfolio separation for elliptical distributions follows Owen and Rabinovitch (1983).

Our results sit most naturally in the elliptical class, which is the broadest family for which the full portfolio theory apparatus (mean-variance optimality, two-fund separation) is known to hold.

Summary. Table 9 classifies the assumptions by their role in each result.

Table 9: Assumptions required for each result.

Result	A1 (moments)	A2 (loc-scale)	A3 (prop. tax)	A4 (part. equil.)
Prop 1′: CV invariance	Required	Not required	**Essential**	Not required
Prop 2′a: Tangency inv.	Required	Not required	**Essential**	Not required
Prop 2′b: Full weight inv.	Required	MV only	**Essential**	Not required
Prop 3′: Orthogonality	Required	Required	**Essential**	Not required
Prop 4′: Pricing neut.	Not req.	Not required	**Essential**	**Essential**

"MV only" means A2 is needed for mean-variance preferences but not for CRRA.

The single assumption that is essential to all four results is **proportionality**—the tax is a fixed fraction of market value, applied uniformly to all assets. This is what generates the multiplicative separability $W_n^A = (1 - \tau_w)^n W_n^B$, which is the structural foundation for everything that follows.

The distributional assumption, by contrast, is a regularity condition that ensures the relevant moments and geometric objects are well-defined. It plays no substantive economic role.

C Continuous Time vs. Discrete Time

The continuous-time GBM formulation (Sections 3–5) and the discrete-time generalisation (Section 7) yield the same qualitative results but differ in detail:

Table 10: Continuous-time vs. discrete-time comparison.

Feature	Continuous time (GBM)	Discrete time (general)
Tax on drift	Additive: $\mu_W = \mu_P - \tau_w$	Multiplicative: $\mu_W = (1-\tau_w)(1+\mu_P) - 1$
Tax on volatility	None: $\sigma_W = \sigma_P$	Scaling: $\sigma_W = (1-\tau_w)\sigma_P$
Sharpe ratio	Invariant	Invariant
Tangency portfolio	Invariant	Invariant
Full weight vector	Invariant for all $f(\mu_W, \sigma_W)$	CRRA: invariant; MV: scales by $\frac{1}{1-\tau_w}$
Geometry in (σ, μ)	Vertical translation	Homothetic contraction

The continuous-time results are the natural limit of the discrete-time results as the period length shrinks.

D Relation to the Existing Literature

Bjerksund and Schjelderup (2022) derive Proposition 8 (pricing neutrality) in a discrete-time DCF framework without specifying a return distribution. Their result that $\mathrm{NPV}_0^D = \mathrm{NPV}_0^F = c_0/r$ is the pricing counterpart of our framework, and rests on no-arbitrage rather than on any specific asset pricing model (see Section 8.2). Our contribution is to embed this in a unified treatment that also covers the portfolio choice and risk-reward dimensions. A technical difference concerns the after-tax discount rate. Bjerksund and Schjelderup assume a beginning-of-period tax base, for which the after-tax cost of capital is $k - \tau_w$ (Kruschwitz et al., 2023). Our model assesses the tax at period end on current market values, yielding $k^A = (1-\tau_w)(1+k) - 1 = k - \tau_w - \tau_w k$ (Equation (57)). The two expressions differ by the cross term $\tau_w k$, which is small for realistic parameters ($\tau_w = 1\%$, $k = 8\%$ gives $\tau_w k = 0.08\%$) but is needed for no-arbitrage under end-of-period assessment. In both cases, the pricing neutrality result $V = V^0$ obtains because the tax reduces the numerator and denominator of the valuation in the same proportion; the timing convention determines only the form of the discount rate, not the conclusion.

Eikseth and Lindset (2009) develop a CAPM-like framework with heterogeneous asset taxes across investors. They derive an after-tax beta equal to the pre-tax beta multiplied by an asset-specific tax adjustment, and show that the Security Market Line becomes a Security Market Fan. Their setting is more general than ours (different tax rates on different assets), but in the special case of a uniform proportional tax, their results are consistent with the portfolio invariance and orthogonality we derive here. Section 8.3.5 shows that book-value taxation (Section 9.2) creates exactly the heterogeneous effective rates that produce the Security Market Fan, providing the equilibrium mechanism behind the beta-dependent GE pricing formula in Proposition 9(b).

Fama (2021) argues that wealth taxes lower asset prices by raising the pre-tax required return: in a Gordon perpetuity, $P = D/(r + \tau_w) < D/r$. His analysis adds the wealth tax to the cost of capital but does not adjust the discount rate for the fact that the entire opportunity set—including the risk-free asset—is also taxed. Section 8.3.4 shows that when the wealth tax applies universally, the correct after-tax discount rate is $k^A = (1 - \tau_w)(1 + k) - 1 < k$, and the price effect vanishes ($V = V^0$). Fama's result holds only when the marginal investor has access to an untaxed outside option, making it a partial-equilibrium statement despite its general-equilibrium framing.

Johnsen and Lensberg (2014) model the wealth tax as a separate claim on the firm using a Modigliani-Miller framework. For listed firms taxed on market value, they derive $V = V^0(1 - \tau_w/k)$, where k is the pre-tax cost of capital. As discussed in Section 8.4, this formula reduces the cash flows by the tax but retains the pre-tax discount rate; using the after-tax discount rate recovers $V = V^0$, consistent with our pricing neutrality result. Their analysis of non-listed firms, where the tax base is book value rather than market value, identifies a leverage-like increase in equity beta that is specific to the Norwegian institutional setting and lies outside our framework of proportional taxation on market value.

Stowe (2021) analyses the capitalisation of wealth taxes into stock and bond valuations, showing that taxed assets will have lower valuations and higher required rates of return. His focus on valuation impacts across asset classes complements our partial equilibrium analysis of pricing neutrality between taxed and untaxed investors.

Hamada and Valdez (2008) prove that the CAPM holds when returns follow multivariate elliptical distributions, using a generalised Stein's lemma. Their result shows that the elliptical class—within which our Propositions $2'$ and $3'$ sit most naturally—is compatible with the full portfolio theory apparatus. As discussed in Section 8.2, our results do not *require* CAPM; the Hamada–Valdez result establishes that working within the elliptical class is consistent with it.

Sandmo (1985) and Stiglitz (1969) analyse the effects of taxation on risk-taking in a broader setting that includes income taxes, capital gains taxes, and their interaction with portfolio choice. The orthogonality result (Proposition 7) is specific to the proportional wealth tax—income taxes operate on returns, not on the quantity of assets, and therefore interact with portfolio choice in fundamentally different ways.

The GBM framework and optimal portfolio choice follow Merton (1969, 1971). The mean-variance portfolio theory follows Markowitz (1952). The interaction of taxation with portfolio choice has been studied by, among others, Constantinides (1983) and Dammon et al. (2001), though the specific results on proportional wealth taxation and orthogonality presented here do not appear to be widely documented in the existing literature.

On the empirical side, Bjørneby et al. (2023) find a positive causal relationship between wealth tax liability and employment in Norwegian closely held firms. Their proposed mechanism—that intangible assets in non-traded firms are effectively tax-exempt, incentivising owners to invest in human capital within their businesses rather than in taxed financial assets—is consistent with the non-uniform assessment channel discussed in Section 9.2 and provides empirical support for the neutrality benchmark: when effective tax rates differ across assets, investors reallocate toward the less-taxed margin rather than reducing economic activity.