

Flow Taxes, Stock Taxes, and Portfolio Choice: A Generalised Neutrality Result

Anders G Frøseth*

March 2026

Abstract

A proportional wealth tax—a levy on the stock of wealth—preserves portfolio neutrality by acting as a uniform drift shift in the Fokker–Planck equation for wealth dynamics. We extend this result to the full system of ownership taxes (*eierkostnader*) that a shareholder faces: a corporate tax on gross profits, a capital income tax on the risk-free return, a dividend and capital gains tax on the excess return, and a wealth tax on net assets. Each tax modifies the drift of the wealth process in a distinct way—multiplicative rescaling, constant shift, or regime-dependent compression—while leaving the diffusion coefficient unchanged. We show that the combined system preserves portfolio neutrality under three conditions: (i) the capital income tax rate equals the corporate tax rate, (ii) the shielding rate equals the risk-free rate, and (iii) the wealth tax assessment is uniform across assets. When these conditions hold, the after-tax excess return is a uniform rescaling of the pre-tax excess return by the factor $(1 - \tau_c)(1 - \tau_d)$, and the drift-shift symmetry of the wealth-tax-only case generalises to a *drift-shift-and-rescale symmetry*. We classify the distortions that arise when each condition fails and show that flow-tax distortions and stock-tax distortions are additively separable: they do not interact. The shielding deduction—a feature of several real-world tax systems, including the Norwegian aksjonærmodellen—emerges as the mechanism that restores the symmetry between equity and debt taxation within this framework. Calibrated to the Norwegian dual income tax, conditions (i) and (ii) hold by institutional design; the only binding distortion is non-uniform wealth tax assessment, which generates portfolio tilts roughly 300 times larger than any residual flow-tax channel.

JEL Classification: G11, G12, H21, H24, H25.

Keywords: Wealth tax, ownership costs, portfolio choice, Fokker–Planck equation, drift-shift symmetry, tax neutrality, shielding deduction, corporate tax, dividend tax.

Contents

1 Introduction

4

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

2	The Fokker–Planck Framework	5
2.1	Wealth dynamics and the Fokker–Planck equation	6
2.2	The drift-shift symmetry	6
2.3	Notation	6
3	Flow Taxes as Drift Modifications	7
3.1	Corporate tax	7
3.2	Capital income tax	8
3.3	Dividend and capital gains tax	8
3.4	Combined flow-tax drift	8
3.4.1	After-tax return on equity asset i	8
3.4.2	After-tax return on the risk-free asset	9
3.4.3	After-tax excess return	9
4	The Generalised Neutrality Theorem	9
4.1	Simplification under (C1) and (C2)	9
4.2	Including the wealth tax	10
4.3	The drift-shift-and-rescale transformation	10
4.4	Main result	10
5	Symmetry Breaking	12
5.1	Violation of (C1): $\tau_k \neq \tau_c$	13
5.2	Violation of (C2): $r_s \neq r_f$	13
5.3	Violation of (C3): $\alpha_i \neq \alpha_j$	14
5.4	Summary	14
6	Flow Taxes and the Cost of Wealth Tax Payment	14
6.1	Dividend extraction	15
6.2	Share sales and liquidity frictions	15
6.3	Indirect pricing effects	16
6.4	Migration	16
6.5	Summary	16
7	The Shielding Mechanism	16
7.1	Definition and mechanics	16
7.2	Symmetry restoration	17
7.3	The shielding deduction in the FP framework	17
8	Interaction Between Flow and Stock Taxes	18
8.1	Additive separability	18
8.2	Drift decomposition	19
8.3	Implications for the wealth distribution	19
9	Results in Mean–Variance Framework	19
9.1	Opportunity set and Sharpe ratios	19

9.2	Portfolio distortions under symmetry breaking	20
9.3	Reduction to the wealth-tax-only case	20
10	The Norwegian Tax System	20
10.1	The dual income tax and conditions (C1)–(C2)	20
10.2	Valuation discounts and condition (C3)	21
10.3	Distortion magnitudes under current parameters	22
10.4	Beyond the proportional framework	23
11	Discussion	25
11.1	Connection to the paper series	25
11.2	Equilibrium prices under inelastic markets	26
11.3	The role of the shielding deduction	26
11.4	Policy implications	27
11.5	Limitations	27
12	Conclusion	27

1 Introduction

The question of whether taxation distorts portfolio composition has a long lineage. [Domar and Musgrave \(1944\)](#) show that a proportional income tax with full loss offset preserves the expected return per unit of risk, so that the investor’s optimal risk exposure is invariant to the tax rate. [Stiglitz \(1969\)](#) extends the result to capital gains and wealth taxation under specific conditions, and [Sandmo \(1977\)](#) provides a general comparative-statics treatment with many assets.

The portfolio neutrality of a proportional wealth tax specifically is now well established. [Frøseth \(2026b\)](#) shows that a uniform levy on the market value of all assets leaves portfolio weights, Sharpe ratios, and asset prices unchanged. [Frøseth \(2026a\)](#) extends the result to stochastic volatility, Epstein–Zin preferences, and general Markov diffusions, while identifying non-uniform assessment as the principal source of distortion in practice. [Frøseth \(2026d\)](#) recasts the neutrality result as a drift-shift symmetry in the Fokker–Planck equation for wealth dynamics: the proportional wealth tax shifts the drift velocity uniformly, $v \rightarrow v - \tau_w$, without coupling to the state, the diffusion coefficient, or the portfolio composition.

These results are derived under the simplifying assumption that the wealth tax is the *only* tax. In practice, no shareholder faces a wealth tax in isolation. A Norwegian personal shareholder, for example, faces four layers of taxation on equity ownership—the *eierkostnader* (ownership costs):

1. A *corporate tax* at rate τ_c on company profits, reducing the gross return available for distribution.
2. A *capital income tax* at rate τ_k on the risk-free component of the return (the portion covered by the shielding deduction).
3. A *dividend and capital gains tax* at rate τ_d on the excess return above the shielding rate.
4. A *wealth tax* at rate τ_w on net assets (with asset-class-specific valuation discounts).

The first three are taxes on the *flow* (the return on wealth); the fourth is a tax on the *stock* (the level of wealth). The distinction matters for the Fokker–Planck framework: flow taxes modify the drift through the return channel, while the stock tax modifies the drift through a direct levy on the state variable.

This paper asks: does portfolio neutrality survive when all four taxes are present? The answer is yes, under three conditions that turn out to be economically natural and, in the Norwegian case, satisfied by institutional design. The conditions are:

- (C1) The capital income tax rate equals the corporate tax rate: $\tau_k = \tau_c$.
- (C2) The shielding rate equals the risk-free rate: $r_s = r_f$.
- (C3) The wealth tax assessment is uniform across assets: $\alpha_i = \alpha$ for all i .

Under (C1)–(C3), the combined tax system acts on the Fokker–Planck equation through two modifications: a uniform drift shift (from the wealth tax) and a uniform rescaling of excess drift

velocities by the factor $(1 - \tau_c)(1 - \tau_d)$ (from the flow taxes). Neither modification alters the *relative* drifts between assets. The drift-shift symmetry of Frøseth (2026d) generalises to a *drift-shift-and-rescale symmetry*, and all the neutrality consequences—invariant portfolio weights, preserved Sharpe ratios, undistorted asset prices—carry through.

When any of (C1)–(C3) fails, the symmetry breaks in a specific and classifiable way. We show that:

- Violating (C1) or (C2) distorts the equity–debt split but preserves the tangency portfolio among risky assets.
- Violating (C3)—non-uniform wealth tax assessment—distorts the tangency portfolio itself, exactly as in Frøseth (2026a).
- The distortions from flow taxes and from the wealth tax are *additively separable*: they do not interact. The presence of a wealth tax neither amplifies nor dampens any flow-tax distortion, and vice versa.

The shielding deduction (Norwegian: *skjermingsfradrag*) emerges as the institutional mechanism that enforces condition (C2). By exempting the risk-free component of the equity return from the elevated dividend tax rate, it ensures that the boundary between the two personal tax regimes aligns with the economic risk-free rate. In the language of the Fokker–Planck framework, the shielding deduction is a *symmetry-restoring device*: it prevents the two-regime tax structure from breaking the uniform rescaling of excess drifts.

The paper is organised as follows. Section 2 recapitulates the Fokker–Planck framework for wealth dynamics and the drift-shift symmetry from Frøseth (2026d). Section 3 introduces the three flow taxes and derives how each modifies the drift of the wealth process. Section 4 states and proves the generalised neutrality theorem. Section 5 classifies the symmetry-breaking channels when each condition fails. Section 6 examines how flow taxes amplify the cost of wealth tax payment through the channels identified in Papers 1 and 3. Section 7 analyses the shielding deduction as a symmetry restoration mechanism. Section 8 establishes the additive separability of flow and stock tax distortions. Section 9 restates the main results in the mean–variance framework of Frøseth (2026b) and Frøseth (2026a). Section 10 calibrates the framework to the Norwegian dual income tax, evaluating conditions (C1)–(C3) against the institutional design and quantifying distortion magnitudes. Section 11 discusses the economic interpretation, the connection to Papers 1–4, and the implications for tax policy design. Section 12 concludes.

2 The Fokker–Planck Framework

This section recapitulates the elements of the Fokker–Planck framework developed in Frøseth (2026d)—itself inspired by the wealth-dynamics model of Bouchaud and Mézard (2000)—that are needed for the present analysis.

2.1 Wealth dynamics and the Fokker–Planck equation

An investor holds a portfolio of risky assets and a risk-free asset. Under geometric Brownian motion, the investor’s wealth evolves as

$$\frac{dW}{W} = \left[r_f + \mathbf{w}^\top (\boldsymbol{\mu} - r_f \mathbf{1}) \right] dt + \mathbf{w}^\top \boldsymbol{\Sigma} d\mathbf{Z}, \quad (1)$$

where $\boldsymbol{\mu}$ is the vector of expected returns, $\boldsymbol{\Sigma}$ is the volatility matrix, $\mathbf{V} = \boldsymbol{\Sigma} \boldsymbol{\Sigma}^\top$ is the covariance matrix, r_f is the risk-free rate, and \mathbf{w} is the vector of portfolio weights.

For a given portfolio, the dynamics of log-wealth $x = \ln W$ are

$$dx = v dt + \sigma_P dB, \quad (2)$$

where $v = r_f + \mathbf{w}^\top (\boldsymbol{\mu} - r_f \mathbf{1}) - \sigma_P^2/2$ is the drift velocity, $\sigma_P^2 = \mathbf{w}^\top \mathbf{V} \mathbf{w}$ is the portfolio variance, and B is a standard Brownian motion.

The Fokker–Planck equation for the density $\pi(x, t)$ of log-wealth across an ensemble of investors is

$$\frac{\partial \pi}{\partial t} = -v \frac{\partial \pi}{\partial x} + D \frac{\partial^2 \pi}{\partial x^2}, \quad (3)$$

with diffusion coefficient $D = \sigma_P^2/2$.

2.2 The drift-shift symmetry

Frøseth (2026d) defines the drift-shift transformation and establishes neutrality as an invariance property.

Definition 1 (Drift-shift transformation, Frøseth (2026d, Definition 1)). For $\tau_w \geq 0$, define the map $\mathcal{T}_\tau : v \mapsto v - \tau_w, D \mapsto D$. The taxed Fokker–Planck operator is $\mathcal{L}_\tau = \mathcal{T}_\tau \circ \mathcal{L}_0$.

Proposition 1 (Neutrality as invariance, Frøseth (2026d, Proposition 2)). *Let assets i and j have drift–diffusion parameters (v_i, D_i) and (v_j, D_j) . Under a proportional wealth tax at rate τ_w :*

1. *The difference in drift velocities is unchanged: $(v_i - \tau_w) - (v_j - \tau_w) = v_i - v_j$.*
2. *The Sharpe-ratio-like quantity $(v_i - v_j)/\sqrt{D_i + D_j - 2D_{ij}}$ is unchanged for all pairs.*
3. *The optimal portfolio weights are unchanged.*

The key property is that the wealth tax shifts all drifts by the *same* constant. Relative drifts, which determine portfolio choice, are invariant. In physical terms, the tax is a uniform external field that does not couple to the internal degrees of freedom of the system.

2.3 Notation

Table 1 collects the symbols used throughout the paper. The first group is inherited from Papers 1–3; the second group contains the flow-tax parameters introduced here.

Table 1: Notation reference. “Paper” indicates where the symbol first appears in the series: Paper 1 = Frøseth (2026b), Paper 2 = Frøseth (2026a), Paper 3 = Frøseth (2026d), Paper 5 = the present paper.

Symbol	Definition	Paper
<i>Portfolio and return primitives</i>		
$\boldsymbol{\mu} = (\mu_1, \dots, \mu_K)^\top$	Vector of gross expected returns	1
$\mathbf{V} = \boldsymbol{\Sigma}\boldsymbol{\Sigma}^\top$	Return covariance matrix	1
r_f	Risk-free rate	1
$\mathbf{w} = (w_1, \dots, w_K)^\top$	Portfolio weights on risky assets	1
γ	Relative risk aversion (CRRA)	1
$\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_K)^\top$	Assessment fractions (wealth tax base)	2
<i>Tax parameters</i>		
τ_c	Corporate tax rate	5
τ_k	Capital income tax rate	5
τ_d	Dividend and capital gains tax rate	5
r_s	Shielding rate (skjermingsrente)	5
τ_w	Wealth tax rate	1
<i>Fokker–Planck framework</i>		
$x = \ln W$	Log-wealth	3
v	Drift velocity of log-wealth	3
$D = \sigma_P^2/2$	Diffusion coefficient	3
$\pi(x, t)$	Density of log-wealth across investors	3
\mathcal{T}_τ	Drift-shift transformation	3
\mathcal{T}^{gen}	Drift-shift-and-rescale transformation	5
<i>Derived quantities</i>		
R_i	After-tax return on equity asset i	5
R_0	After-tax return on risk-free asset	5
$R_i^{\text{ex}} = R_i - R_0$	After-tax excess return	5

3 Flow Taxes as Drift Modifications

We now introduce three flow taxes and derive how each modifies the drift of the wealth process. Throughout, we consider K risky assets held in corporate form and one risk-free asset held personally (e.g. bank deposits).

3.1 Corporate tax

A corporate tax at rate $\tau_c \in (0, 1)$ is levied on company profits. The gross return μ_i on asset i held in corporate form becomes $\mu_i(1 - \tau_c)$ after corporate tax. The risk-free rate r_f , earned on assets held personally, is not subject to corporate tax.

In log-wealth coordinates, the drift contribution of asset i changes from μ_i to $\mu_i(1 - \tau_c)$. The corporate tax is a *multiplicative* modification of the gross return:

$$v_i \rightarrow v_i^{(c)} = (1 - \tau_c)\mu_i - \frac{1}{2}\sigma_i^2. \quad (4)$$

Remark (Asymmetry between equity and debt). The corporate tax applies to company profits but not to the risk-free return on bank deposits. This creates an asymmetry: the after-tax return on equity is $\mu_i(1 - \tau_c)$, while the after-tax return on deposits is r_f . The excess return of equity over deposits becomes $\mu_i(1 - \tau_c) - r_f$, which is *not* a uniform scaling of the pre-tax excess return $\mu_i - r_f$. The corporate tax alone therefore distorts the equity–debt allocation. This asymmetry is resolved when the capital income tax is added (Section 3.4).

3.2 Capital income tax

A capital income tax at rate τ_k is levied on the risk-free return component. For bank deposits, the full return r_f is taxed:

$$r_f \rightarrow r_f(1 - \tau_k). \quad (5)$$

For equities, the capital income tax applies to the *shielded* portion of the return—the part up to the shielding rate r_s —at rate τ_k . The shielding mechanism is described in Section 7; for now, we work with the generic rate τ_k on the risk-free component.

3.3 Dividend and capital gains tax

A dividend and capital gains tax at rate τ_d is levied on the *excess* of the after-corporate-tax return over the shielding rate r_s . For an equity asset with after-corporate-tax return $\mu_i(1 - \tau_c) > r_s$:

$$\text{Tax} = \tau_d \cdot [\mu_i(1 - \tau_c) - r_s]. \quad (6)$$

The dividend tax is a *multiplicative* modification of the excess return: it compresses the portion of the return above r_s by the factor $(1 - \tau_d)$.

Remark (Dividend–gains symmetry). In the Norwegian system, dividends and capital gains are taxed at the same effective rate (the nominal 22% multiplied by the upward adjustment factor of 1.72, giving 37.84%). Losses are deductible at the same rate. This symmetry ensures that the choice between dividends and capital gains as a source of cash to pay the wealth tax is not distorted by the flow taxes. In the Fokker–Planck framework, only the *total* after-tax return matters for the drift; the composition between dividends and realisations is irrelevant.

3.4 Combined flow-tax drift

We now derive the drift of the wealth process under all three flow taxes.

3.4.1 After-tax return on equity asset i

Assume $\mu_i(1 - \tau_c) > r_s$. The after-tax return on equity asset i , after corporate tax, capital income tax on the shielded portion, and dividend tax on the excess, is:

$$R_i = \mu_i(1 - \tau_c) - \tau_k \cdot r_s - \tau_d[\mu_i(1 - \tau_c) - r_s]. \quad (7)$$

Collecting terms:

$$R_i = \mu_i(1 - \tau_c)(1 - \tau_d) + r_s(\tau_d - \tau_k). \quad (8)$$

3.4.2 After-tax return on the risk-free asset

$$R_0 = r_f(1 - \tau_k). \quad (9)$$

3.4.3 After-tax excess return

The after-tax excess return of equity asset i over the risk-free asset is:

$$R_i^{\text{ex}} = R_i - R_0 = \mu_i(1 - \tau_c)(1 - \tau_d) + r_s(\tau_d - \tau_k) - r_f(1 - \tau_k). \quad (10)$$

This is the general expression. We now examine its structure under the conditions (C1) and (C2).

Table 2 summarises the drift modification introduced by each tax layer.

Table 2: Drift modification by tax layer. Each row shows how a single tax modifies the drift velocity v_i of asset i in log-wealth coordinates. The ‘‘Character’’ column classifies the modification by its dependence on the asset index.

Tax	Rate	Drift modification	Character
Corporate	τ_c	$\mu_i \rightarrow \mu_i(1 - \tau_c)$	Multiplicative on gross
Capital income	τ_k	$r_f \rightarrow r_f(1 - \tau_k)$	Shift on risk-free drift
Dividend/gains	τ_d	$(1 - \tau_d)[\mu_i(1 - \tau_c) - r_s]$	Multiplicative on excess
Wealth	τ_w	$v_i \rightarrow v_i - \tau_w \alpha_i$	Uniform shift (if $\alpha_i = \alpha$)

4 The Generalised Neutrality Theorem

4.1 Simplification under (C1) and (C2)

Condition (C1): $\tau_k = \tau_c$. The capital income tax rate equals the corporate tax rate.

Condition (C2): $r_s = r_f$. The shielding rate equals the risk-free rate.

Under (C1) and (C2), the constant terms in (10) simplify. Substituting $\tau_k = \tau_c$ and $r_s = r_f$:

$$\begin{aligned} R_i^{\text{ex}} &= \mu_i(1 - \tau_c)(1 - \tau_d) + r_f(\tau_d - \tau_c) - r_f(1 - \tau_c) \\ &= \mu_i(1 - \tau_c)(1 - \tau_d) + r_f\tau_d - r_f\tau_c - r_f + r_f\tau_c \\ &= \mu_i(1 - \tau_c)(1 - \tau_d) - r_f(1 - \tau_d) \\ &= (1 - \tau_d)[\mu_i(1 - \tau_c) - r_f]. \end{aligned} \quad (11)$$

Now observe that $\mu_i(1 - \tau_c) - r_f = (1 - \tau_c)\mu_i - r_f$. For the *difference* between two risky assets i and j :

$$R_i^{\text{ex}} - R_j^{\text{ex}} = (1 - \tau_d)(1 - \tau_c)(\mu_i - \mu_j). \quad (12)$$

This is a *uniform scaling* of the pre-tax return differences. The scalar factor $(1 - \tau_c)(1 - \tau_d)$ is common to all asset pairs.

4.2 Including the wealth tax

Adding the wealth tax at rate τ_w with assessment fractions α_i for risky asset i and α_0 for the risk-free asset:

$$R_i^{\text{full}} = (1 - \tau_d)[\mu_i(1 - \tau_c) - r_f] - \tau_w(\alpha_i - \alpha_0). \quad (13)$$

Under **condition (C3)**, $\alpha_i = \alpha$ for all i , the wealth tax term cancels in excess returns between risky assets:

$$R_i^{\text{full}} - R_j^{\text{full}} = (1 - \tau_d)(1 - \tau_c)(\mu_i - \mu_j). \quad (14)$$

4.3 The drift-shift-and-rescale transformation

Under (C1)–(C3), the combined tax system acts on the drift of the wealth process through two modifications:

1. A *uniform rescaling* of excess drift velocities by $(1 - \tau_c)(1 - \tau_d)$, from the flow taxes.
2. A *uniform shift* of all drifts by $-\tau_w\alpha$, from the wealth tax.

Definition 2 (Drift-shift-and-rescale transformation). For tax parameters $(\tau_c, \tau_d, \tau_w, \alpha)$ satisfying (C1)–(C3), define the map

$$\mathcal{T}^{\text{gen}} : \begin{cases} v_i - v_0 \mapsto (1 - \tau_c)(1 - \tau_d)(v_i - v_0) \\ D_i \mapsto D_i \end{cases} \quad (15)$$

where v_0 is the risk-free drift. The overall drift level is shifted by $-(1 - \tau_d)\tau_c r_f - \tau_w\alpha$.

This generalises Definition 1. The drift-shift transformation of Frøseth (2026d) is the special case $\tau_c = \tau_d = 0$, $\tau_w > 0$.

4.4 Main result

Theorem 1 (Neutrality under combined taxation). *Let asset returns follow an Itô diffusion with well-defined first and second moments. Let the tax system consist of a corporate tax at rate τ_c , a capital income tax at rate τ_k , a dividend/gains tax at rate τ_d with shielding rate r_s , and a proportional wealth tax at rate τ_w with assessment fraction α .*

If the following conditions hold:

$$(C1) \quad \tau_k = \tau_c,$$

$$(C2) \quad r_s = r_f,$$

$$(C3) \quad \alpha_i = \alpha \text{ for all assets } i,$$

then the optimal portfolio weights \mathbf{w}^ are independent of all tax rates (τ_c, τ_d, τ_w) .*

Proof. Under (C1) and (C2), the after-tax excess return of asset i is $(1 - \tau_d)(1 - \tau_c)(\mu_i - r_f) + c$, where c is a constant independent of i (Equation (11)). Under (C3), the wealth tax contributes a further constant $-\tau_w\alpha$ to all excess returns (Equation (13)).

The vector of after-tax excess returns is therefore

$$\mathbf{R}^{\text{ex}} = (1 - \tau_c)(1 - \tau_d)(\boldsymbol{\mu} - r_f \mathbf{1}) + c' \mathbf{1}, \quad (16)$$

where $c' = -(1 - \tau_d)\tau_c r_f - \tau_w(\alpha - \alpha_0)$ is a scalar.

The optimal portfolio in the Markowitz problem satisfies the first-order condition:

$$\mathbf{w}^* = \frac{1}{\gamma} \mathbf{V}^{-1} \mathbf{R}^{\text{ex}}. \quad (17)$$

Substituting (16):

$$\mathbf{w}^* = \frac{(1 - \tau_c)(1 - \tau_d)}{\gamma} \mathbf{V}^{-1}(\boldsymbol{\mu} - r_f \mathbf{1}) + \frac{c'}{\gamma} \mathbf{V}^{-1} \mathbf{1}. \quad (18)$$

The first term is the untaxed tangency portfolio direction, scaled by $(1 - \tau_c)(1 - \tau_d)$. The second term is proportional to $\mathbf{V}^{-1} \mathbf{1}$, the global minimum variance portfolio direction.

Under CRRA preferences in continuous time, the portfolio problem takes the form of the Merton problem. The value function $J(W, t) = W^{1-\gamma} f(t)/(1 - \gamma)$ is homogeneous of degree $1 - \gamma$ in wealth. The first-order condition for the optimal portfolio weight yields

$$\mathbf{w}^* = \frac{1}{\gamma} \mathbf{V}^{-1}(\boldsymbol{\mu}^{\text{after}} - R_0 \mathbf{1}), \quad (19)$$

where $\boldsymbol{\mu}^{\text{after}}$ is the vector of after-tax expected returns and R_0 is the after-tax risk-free rate.

Substituting $\boldsymbol{\mu}^{\text{after}} = (1 - \tau_c)(1 - \tau_d)\boldsymbol{\mu} + [r_s(\tau_d - \tau_k) - \tau_w\alpha]\mathbf{1}$ and $R_0 = r_f(1 - \tau_k) - \tau_w\alpha_0$:

$$\begin{aligned} \mathbf{w}^* &= \frac{1}{\gamma} \mathbf{V}^{-1}[(1 - \tau_c)(1 - \tau_d)\boldsymbol{\mu} + [r_s(\tau_d - \tau_k) - \tau_w\alpha]\mathbf{1} - [r_f(1 - \tau_k) - \tau_w\alpha_0]\mathbf{1}] \\ &= \frac{(1 - \tau_c)(1 - \tau_d)}{\gamma} \mathbf{V}^{-1}(\boldsymbol{\mu} - r_f \mathbf{1}), \end{aligned} \quad (20)$$

where the last equality uses (C1), (C2), and (C3) to cancel all constant terms (as computed in Equation (11) and Equation (14)).

The optimal weight \mathbf{w}^* in (20) is proportional to $\mathbf{V}^{-1}(\boldsymbol{\mu} - r_f \mathbf{1})$ —the same direction as the untaxed portfolio. The scalar $(1 - \tau_c)(1 - \tau_d)/\gamma$ determines the total allocation to risky assets but not the composition. Under CRRA, the total allocation is also invariant because the homogeneity of the value function absorbs the scaling factor into the effective rate of time preference (see Frøseth (2026a), Proposition 1, for the detailed argument under stochastic volatility).

Since neither the direction nor the magnitude of \mathbf{w}^* depends on τ_c , τ_d , or τ_w , portfolio neutrality holds under the combined tax system. \square

Remark (Distribution-free extension). The drift-shift-and-rescale transformation does not depend on the Gaussian assumption. By the same argument as in Frøseth (2026d), Proposition 3 (distribution-free drift shift), the result extends to general Itô diffusions with state-dependent drift and volatility, provided preferences are CRRA and the return dynamics are

wealth-independent. The flow-tax rescaling factor $(1 - \tau_c)(1 - \tau_d)$ multiplies the excess drift regardless of the distributional form of the noise.

Figure 1 illustrates the theorem in mean–standard deviation space. The combined tax system acts in two stages: flow taxes rescale excess returns by $(1 - \tau_c)(1 - \tau_d)$, contracting the efficient frontier and the capital allocation line toward the risk-free rate; the wealth tax then shifts the entire opportunity set vertically by $-\tau_w\alpha$. At each stage the tangency portfolio remains at the same volatility σ^* , and the Sharpe ratio is preserved.

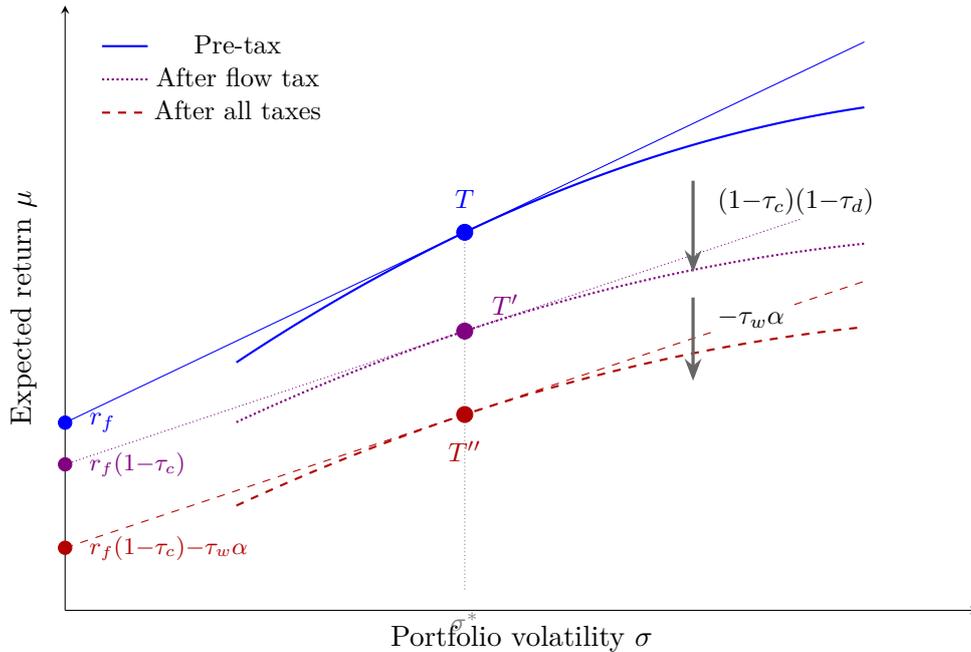


Figure 1: Generalised neutrality in mean–standard deviation space. Flow taxes reduce the risk-free rate from r_f to $r_f(1 - \tau_c)$ and rescale excess returns by $(1 - \tau_c)(1 - \tau_d)$, contracting the efficient frontier and capital allocation line. The wealth tax then shifts the entire opportunity set vertically by $-\tau_w\alpha$. At both stages the tangency portfolio remains at the same volatility σ^* : the combined tax system is neutral with respect to portfolio composition. This generalises the vertical-translation result of Frøseth (2026b, Proposition 3).

Remark (Why the constant term vanishes). The constant $c'\mathbf{1}$ in (16) shifts all excess returns by the same amount. This is the flow-tax analogue of the wealth tax’s uniform drift shift. In the Markowitz problem, this constant affects the total allocation to risky assets (the distance along the capital allocation line) but not the tangency portfolio direction. Under CRRA in continuous time, even the total allocation is invariant, because the constant enters the effective discount rate and is absorbed by the homogeneity of the value function.

5 Symmetry Breaking

When any of (C1)–(C3) fails, the drift-shift-and-rescale symmetry breaks. This section classifies the distortions by the type of modification they introduce.

5.1 Violation of (C1): $\tau_k \neq \tau_c$

When the capital income tax rate differs from the corporate tax rate, the first-level taxation of equity and debt is asymmetric. The after-tax return on equity is scaled by $(1 - \tau_c)$, while the after-tax return on deposits is scaled by $(1 - \tau_k)$.

The after-tax excess return of risky asset i over the risk-free asset becomes (from Equation (10)):

$$R_i^{\text{ex}} = (1 - \tau_d)[\mu_i(1 - \tau_c) - r_f] + r_f(\tau_d - \tau_k) - r_f(1 - \tau_k) + r_f(1 - \tau_d) + r_f(\tau_k - \tau_c). \quad (21)$$

The difference between two risky assets is still:

$$R_i^{\text{ex}} - R_j^{\text{ex}} = (1 - \tau_d)(1 - \tau_c)(\mu_i - \mu_j), \quad (22)$$

which is a uniform scaling. **The tangency portfolio among risky assets is preserved.**

The distortion appears in the equity–debt split. The effective after-tax risk-free rate is $r_f(1 - \tau_k)$, while equity returns are scaled by $(1 - \tau_c)$. When $\tau_k > \tau_c$, deposits are penalised relative to equity; when $\tau_k < \tau_c$, the reverse. This is a tilt along the capital allocation line—a change in the total allocation to risky assets—but not a rotation of the tangency portfolio.

FP interpretation: The excess drift between any two risky assets is uniformly rescaled. The asymmetry appears only in the drift of equity relative to the risk-free asset, as a non-uniform constant shift.

5.2 Violation of (C2): $r_s \neq r_f$

When the shielding rate differs from the risk-free rate, the boundary between the two personal tax regimes (τ_k and τ_d) is misaligned with the economic risk-free rate. The after-tax excess return contains the term $r_s(\tau_d - \tau_k) - r_f(1 - \tau_k)$, which under (C1) becomes $r_s(\tau_d - \tau_c) - r_f(1 - \tau_c)$. When $r_s \neq r_f$, this does not simplify to $-r_f(1 - \tau_d)$.

The discrepancy is $(r_s - r_f)(\tau_d - \tau_c)$: a constant that shifts all equity excess returns by the same amount. The effect is the same in character as a violation of (C1)—a tilt along the capital allocation line—since the constant shift does not depend on the asset.

The tangency portfolio is preserved. Only the equity–debt allocation is distorted.

Remark (Practical relevance). In Norway, the shielding rate is based on the arithmetic average of three-month treasury bill (*statskasseveksel*) rates, plus a markup of 0.5 percentage points.¹ By construction, $r_s \geq r_f$ whenever treasury bills are a reasonable proxy for r_f , so the (C2) distortion has a known sign: it favours equity over debt. In either case, the magnitude is small—see Section 10 for the quantitative bound—and of second order relative to the distortion from non-uniform assessment (C3).

¹The markup was introduced by the Skauge Committee (NOU 2003:9, 2003) to compensate for the asymmetry that unused shielding is lost upon realisation. The Torvik Committee (NOU 2022:20, 2022) proposed replacing the current basis with ten-year government bond yields (without the markup), arguing that the neutral shielding rate lies between the risk-free rate and the investor’s marginal financing cost; see §8.5.7.4 of NOU 2022:20 (2022).

5.3 Violation of (C3): $\alpha_i \neq \alpha_j$

This is the result of Frøseth (2026a), §4. When assessment fractions differ across assets—a feature whose distortionary implications for asset valuation are also analysed by Bjerksund and Schjelderup (2022)—the wealth tax creates an asset-specific wedge $-\tau_w(\alpha_i - \alpha_0)$ in the excess return. The portfolio distortion is:

$$\Delta \mathbf{w}^* = -\frac{\tau_w}{\gamma} \mathbf{V}^{-1}(\boldsymbol{\alpha} - \alpha_0 \mathbf{1}), \quad (23)$$

where $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_K)^\top$. This *does* distort the tangency portfolio.

FP interpretation: The drift shift becomes asset-dependent, $v_i \rightarrow v_i - \tau_w \alpha_i$. This is the anisotropic field of Frøseth (2026d), §7.1 (Channel 1: book-value assessment). The uniformity of the drift-shift transformation is broken.

5.4 Summary

Table 3: Symmetry-breaking classification.

Condition violated	Tangency	Equity–debt	FP character
(C1): $\tau_k \neq \tau_c$	Preserved	Distorted	Constant shift in equity–debt drift
(C2): $r_s \neq r_f$	Preserved	Distorted	Constant shift in equity–debt drift
(C3): $\alpha_i \neq \alpha_j$	Distorted	Distorted	Anisotropic drift shift
All hold	Preserved	Preserved	Uniform rescaling + uniform shift

6 Flow Taxes and the Cost of Wealth Tax Payment

The preceding section analyses how violations of (C1)–(C3) distort portfolio choice. A separate question is how flow taxes interact with the *payment mechanism* of the wealth tax.

Under (C1)–(C3), the flow taxes levied on dividends or capital gains used to pay the wealth tax do not themselves break neutrality: the rates τ_c and τ_d are uniform across assets, so the tax cost of generating cash is the same regardless of which asset is sold or which firm distributes earnings. The tangency portfolio is unaffected.

The interaction arises at a different level. Frøseth (2026b), §9.3 establishes a payment hierarchy for the wealth tax: investors cover the liability first from excess liquidity and other income, then through dividend extraction from the firm, and only as a last resort through partial sale of shares. The wealth tax is collected in quarterly advance instalments (*forskuddsskatt*) during the assessment year, spreading the liquidity demand over time. Frøseth (2026d), §7 reformulates these payment mechanisms as distortion channels in the Fokker–Planck framework.

Flow taxes do not create new channels, but they amplify the existing ones by increasing the gross cost of generating cash at each stage of the hierarchy.

6.1 Dividend extraction

The most common payment mechanism for owners of private firms is dividend extraction (Frøseth 2026b, §9.4). When firm owners extract dividends to pay the wealth tax, those dividends are taxed twice: at the corporate rate τ_c and then at the personal rate τ_d (above the shielding deduction). To deliver $\tau_w W$ to the investor, the firm must generate gross earnings of

$$E = \frac{\tau_w W}{(1 - \tau_c)(1 - \tau_d)} = \frac{\tau_w W}{k}, \quad (24)$$

where $k = (1 - \tau_c)(1 - \tau_d)$ is the combined flow-tax pass-through factor. Under Norwegian rates ($\tau_c = 0.22$, $\tau_d = 0.378$), $k \approx 0.484$, so each krone of wealth tax paid requires approximately two kroner of pre-tax earnings.

This amplifies Channel 3 of Frøseth (2026d): the drain on firm capital becomes

$$\frac{dK}{dt} = f(K) - \delta K - \frac{\tau_w W}{k}, \quad (25)$$

rather than $f(K) - \delta K - \tau_w W$ in the wealth-tax-only case. The coupling between investor wealth and firm capital is strengthened by the factor $1/k$. For firms facing financing frictions, this may force additional equity issuance or reduce investment, lowering the expected return and further reducing the drift.

6.2 Share sales and liquidity frictions

When excess liquidity and dividend capacity are insufficient to cover the liability, the investor must sell shares. Any realised capital gain is taxed at rate τ_d . To generate a net payment of $\tau_w W$, the investor must sell an amount

$$S = \frac{\tau_w W}{1 - \tau_d \cdot g}, \quad (26)$$

where g is the fraction of the sale proceeds that constitutes a taxable gain. When $g > 0$, the required gross sale exceeds the tax liability.

In the Fokker–Planck framework, this amplifies the friction cost in Channel 2 of Frøseth (2026d):

$$c(W, \ell) \rightarrow c\left(\frac{W}{1 - \tau_d g}, \ell\right), \quad (27)$$

increasing the state-dependent drift modification. The effect is largest for investors holding appreciated, illiquid assets—precisely those for whom the original liquidity channel is already most severe.

The quarterly instalment schedule mitigates the liquidity impact by spreading forced sales over the assessment year rather than concentrating them in a single transaction. This limits the instantaneous market impact but does not eliminate the flow-tax amplification of the per-transaction friction cost.

6.3 Indirect pricing effects

The effect on prices is indirect, through the market impact channel (Channel 5 of Frøseth 2026d). The flow-tax amplification of forced sales increases the aggregate selling pressure F . Under the inelastic markets hypothesis (Gabaix and Koijen, 2022), the price depression is proportional to \sqrt{F} . This is an indirect pricing effect of the combined tax system that could not be identified in Papers 1–2, which analysed the wealth tax in isolation: the flow taxes roughly double the gross cash needed at each stage of the payment hierarchy, increasing the equilibrium selling pressure and thus the price impact.

6.4 Migration

The threshold wealth W^* at which emigration becomes attractive depends on the total tax cost, not just the wealth tax. Flow taxes reduce the after-tax return, lowering the effective wealth accumulation rate and thereby shifting W^* downward. The absorbing-boundary mechanism of Frøseth (2026d), §7.4 is unchanged in structure, but the boundary moves.

6.5 Summary

Table 4: Flow-tax amplification of wealth tax payment channels.

Channel	Wealth tax only	With flow taxes	Amplification
Dividends (Ch. 3)	$\tau_w W$ drain	$\tau_w W/k$ drain	Factor $1/k \approx 2.1$ (Norway)
Liquidity (Ch. 2)	$c(W, \ell)$	$c(W/(1 - \tau_d g), \ell)$	$\tau_d g$ on realised gains
Market impact (Ch. 5)	$F(\tau_w)$	$F(\tau_w/(1 - \tau_d g))$	Larger selling pressure
Migration (Ch. 4)	$W^*(\tau_w)$	$W^*(\tau_w, \tau_c, \tau_d)$	Lower threshold

The dividend channel is most severely affected: under Norwegian rates, the amplification factor $1/k \approx 2.1$ means that the effective drain on firm capital is more than double the statutory wealth tax liability. The pricing effect is indirect, operating through the market impact of amplified forced sales rather than through a direct violation of pricing neutrality.

7 The Shielding Mechanism

The shielding deduction (*skjermingsfradrag* in the Norwegian system) is the institutional mechanism that enforces condition (C2). It is a form of rate-of-return allowance—an idea with deep roots in optimal tax design (Boadway and Bruce, 1984; Sørensen, 2005; Mirrlees et al., 2011). This section analyses its role in the neutrality framework.

7.1 Definition and mechanics

The shielding deduction exempts a “normal” return on invested capital from the elevated dividend tax rate. For each share, the shielding amount is:

$$S = (\text{cost basis}) \times r_s, \quad (28)$$

where r_s is the shielding rate (*skjermingsrente*). Institutionally, r_s is set by Skattedirektoratet as the arithmetic average of three-month treasury bill rates plus 0.5 percentage points.² Dividends and capital gains up to the shielding amount are taxed at the capital income rate τ_k ; the excess is taxed at the dividend rate $\tau_d > \tau_k$.

7.2 Symmetry restoration

In the absence of shielding (i.e. the full equity return is taxed at τ_d), the after-tax return on equity asset i would be:

$$R_i^{\text{no shield}} = \mu_i(1 - \tau_c)(1 - \tau_d). \quad (29)$$

The after-tax excess return over the risk-free asset would be:

$$R_i^{\text{no shield}} - R_0 = \mu_i(1 - \tau_c)(1 - \tau_d) - r_f(1 - \tau_k). \quad (30)$$

For two risky assets, the difference is still $(1 - \tau_c)(1 - \tau_d)(\mu_i - \mu_j)$ —a uniform scaling. The tangency portfolio among risky assets is preserved even without shielding.

The distortion from removing shielding appears in the equity–debt split. Without shielding and with $\tau_k = \tau_c$:

$$R_i^{\text{no shield}} - R_0 = (1 - \tau_c)(1 - \tau_d)\mu_i - r_f(1 - \tau_c). \quad (31)$$

This is not a uniform scaling of $\mu_i - r_f$; the risk-free rate is scaled by $(1 - \tau_c)$ while the equity return is scaled by $(1 - \tau_c)(1 - \tau_d)$. The after-tax Sharpe ratio of equity relative to deposits is distorted—equity is penalised by the additional factor $(1 - \tau_d)$ on the *entire* return, not just the excess.

With shielding at rate $r_s = r_f$, the tax on the first r_f of the equity return is at rate $\tau_k = \tau_c$, matching the rate on deposits. Only the excess $\mu_i(1 - \tau_c) - r_f$ is taxed at τ_d . The result, derived in Equation (11), is:

$$R_i - R_0 = (1 - \tau_d)[\mu_i(1 - \tau_c) - r_f], \quad (32)$$

which is a uniform scaling of the after-corporate-tax excess return.

7.3 The shielding deduction in the FP framework

Without shielding, the equity drift is $v_i^{\text{no shield}} = (1 - \tau_c)(1 - \tau_d)\mu_i - \sigma_i^2/2$. The risk-free drift is $v_0 = r_f(1 - \tau_c) - 0$. The excess drift contains the term $(1 - \tau_c)(1 - \tau_d)\mu_i - (1 - \tau_c)r_f = (1 - \tau_c)[(1 - \tau_d)\mu_i - r_f]$, which mixes the scaling factors.

With shielding, the excess drift is $(1 - \tau_d)[(1 - \tau_c)\mu_i - r_f] = (1 - \tau_d)(1 - \tau_c)(\mu_i - r_f) - (1 - \tau_d)\tau_c r_f$. The factor $(1 - \tau_d)(1 - \tau_c)$ applies *uniformly* to the pre-tax excess return $\mu_i - r_f$. The remaining term $-(1 - \tau_d)\tau_c r_f$ is a constant independent of i .

The shielding deduction converts a mixed scaling (equity at $(1 - \tau_c)(1 - \tau_d)$, debt at

²Formally, the published *skjermingsrente* is this pre-tax basis reduced by the alminnelig inntekt rate ($r_s^{\text{AT}} = r_s(1 - \tau_k)$). In the model, r_s denotes the pre-tax rate—the *beregningsgrunnlag*—so that condition (C2) reads $r_s = r_f$ with both rates before personal tax.

$(1 - \tau_c)$) into a uniform scaling of excess returns. This is the symmetry restoration: the drift-shift-and-rescale transformation becomes well-defined only when the shielding aligns the tax regimes at the risk-free boundary.

Remark (Policy interpretation). The shielding deduction is often justified on equity-theoretic grounds: a “normal” return on capital should not face punitive taxation. The Fokker–Planck framework provides a complementary justification from efficiency: the shielding deduction is the mechanism that preserves the drift-shift-and-rescale symmetry, ensuring that the combined tax system does not distort portfolio composition. In this sense, the shielding is not merely equitable—it is the condition for allocative neutrality.

8 Interaction Between Flow and Stock Taxes

Having characterised the neutrality conditions, the payment channels, and the shielding mechanism separately, we now ask how the flow-tax and stock-tax distortions combine.

8.1 Additive separability

Under (C1) and (C2), the after-tax excess return (Equation (13)) has the form:

$$R_i^{\text{full}} = \underbrace{(1 - \tau_d)(1 - \tau_c)(\mu_i - r_f) + \text{const.}}_{\text{flow taxes}} - \underbrace{\tau_w(\alpha_i - \alpha_0)}_{\text{wealth tax}}. \quad (33)$$

The two contributions are additive. There is no cross-term—the flow-tax distortion (if any) and the wealth-tax distortion (from non-uniform assessment) simply add up.

Proposition 2 (Additive separability of distortions). *Under conditions (C1) and (C2), the portfolio distortion from the combined tax system is:*

$$\Delta \mathbf{w}^* = \underbrace{\Delta \mathbf{w}_{\text{flow}}^*}_{\text{zero if (C1)–(C2) hold}} + \underbrace{\Delta \mathbf{w}_{\text{stock}}^*}_{\text{zero if (C3) holds}}. \quad (34)$$

The flow-tax and stock-tax distortions do not interact: the presence of a wealth tax does not amplify or dampen any flow-tax distortion, and vice versa.

Proof. The after-tax excess return vector is $\mathbf{R}^{\text{ex}} = (1 - \tau_c)(1 - \tau_d)(\boldsymbol{\mu} - r_f \mathbf{1}) + c \mathbf{1} - \tau_w(\boldsymbol{\alpha} - \alpha_0 \mathbf{1})$. The optimal portfolio is $\mathbf{w}^* = (1/\gamma)\mathbf{V}^{-1}\mathbf{R}^{\text{ex}}$. The distortion relative to the untaxed case $\mathbf{w}_0^* = (1/\gamma)\mathbf{V}^{-1}(\boldsymbol{\mu} - r_f \mathbf{1})$ is:

$$\begin{aligned} \Delta \mathbf{w}^* &= \frac{1}{\gamma} \mathbf{V}^{-1} [(1 - \tau_c)(1 - \tau_d)(\boldsymbol{\mu} - r_f \mathbf{1}) - (\boldsymbol{\mu} - r_f \mathbf{1})] + \frac{c}{\gamma} \mathbf{V}^{-1} \mathbf{1} - \frac{\tau_w}{\gamma} \mathbf{V}^{-1} (\boldsymbol{\alpha} - \alpha_0 \mathbf{1}) \\ &= \frac{1 - (1 - \tau_c)(1 - \tau_d)}{\gamma} \mathbf{V}^{-1} (\boldsymbol{\mu} - r_f \mathbf{1}) \cdot (-1) + \frac{c}{\gamma} \mathbf{V}^{-1} \mathbf{1} - \frac{\tau_w}{\gamma} \mathbf{V}^{-1} (\boldsymbol{\alpha} - \alpha_0 \mathbf{1}). \end{aligned} \quad (35)$$

Under CRRA in continuous time, the first two terms vanish (absorbed by the homogeneity of the value function, as in the proof of Theorem 1). The remaining distortion is purely from the wealth tax: $\Delta \mathbf{w}^* = -(\tau_w/\gamma)\mathbf{V}^{-1}(\boldsymbol{\alpha} - \alpha_0 \mathbf{1})$, which is independent of τ_c and τ_d . \square

8.2 Drift decomposition

In the Fokker–Planck equation, the combined tax modifies the drift as:

$$v_i^{\text{after}} - v_j^{\text{after}} = (1 - \tau_c)(1 - \tau_d)(v_i - v_j) - \tau_w(\alpha_i - \alpha_j). \quad (36)$$

The first term is a uniform rescaling of pre-tax drift differences (flow taxes); the second is an asset-dependent shift (wealth tax assessment). The two terms are additive and independent.

The flow taxes rescale the drift differences without altering their sign or relative magnitude. The wealth tax assessment introduces an additive distortion that is entirely determined by the assessment fractions α and is independent of the flow-tax rates.

8.3 Implications for the wealth distribution

The Pareto exponent of the stationary wealth distribution under GBM is $\alpha = 1 + 2v/\sigma^2$ (Frøseth 2026d, §3.2). Under the combined tax system, the effective drift is modified by both the flow taxes and the wealth tax. Since the modifications are additive, their effects on the Pareto exponent are also additive:

$$\alpha_{\text{tax}} = 1 + \frac{2[v - \Delta v_{\text{flow}} - \tau_w \alpha]}{\sigma^2}, \quad (37)$$

where Δv_{flow} captures the drift reduction from the flow taxes. The flow taxes steepen the Pareto tail (reduce the drift) independently of the wealth tax, and the effects compound.

9 Results in Mean–Variance Framework

The preceding sections derive the main results in the Fokker–Planck framework of Frøseth (2026d). This section restates them in the mean–variance and portfolio-weight formulation of Frøseth (2026b) and Frøseth (2026a).

9.1 Opportunity set and Sharpe ratios

Corollary 1 (Mean–standard deviation formulation). *Under (C1)–(C3), the combined tax system contracts the mean–standard deviation opportunity set by the factor $(1 - \tau_c)(1 - \tau_d)$ in the mean direction while leaving the standard deviation axis unchanged. In particular:*

1. *The after-tax Sharpe ratio of every portfolio satisfies $\text{SR}^{\text{after}}(p) = (1 - \tau_c)(1 - \tau_d) \text{SR}^{\text{pre}}(p)$. Since the scaling is uniform, the ranking of portfolios by Sharpe ratio is unchanged.*
2. *The tangency portfolio—the portfolio that maximises the Sharpe ratio—is the same before and after tax.*
3. *The capital allocation line contracts toward the after-tax risk-free rate $R_0 = r_f(1 - \tau_c) - \tau_w \alpha$ but retains the same slope direction.*

This generalises the orthogonality result of Frøseth (2026b), Propositions 3 and 7: the wealth tax shifts the opportunity set vertically (drift shift); the flow taxes contract it vertically (drift rescaling). Both operations preserve the tangency portfolio.

9.2 Portfolio distortions under symmetry breaking

Corollary 2 (Distortion classification in portfolio-weight space). *The violations of (C1)–(C3) affect different components of the optimal portfolio:*

1. Violations of (C1) or (C2). *The tangency portfolio direction $\mathbf{V}^{-1}(\boldsymbol{\mu} - r_f \mathbf{1})$ is preserved; only the total allocation to risky assets changes. The distortion is a scalar adjustment to the position along the capital allocation line (Frøseth 2026b, Proposition 3).*
2. Violation of (C3). *The tangency portfolio itself is distorted. The portfolio shift is*

$$\Delta \mathbf{w}^* = -\frac{\tau_w}{\gamma} \mathbf{V}^{-1}(\boldsymbol{\alpha} - \alpha_0 \mathbf{1}),$$

which tilts the portfolio toward assets with lower assessment fractions (Frøseth 2026a, Proposition 5). This is the only violation that alters the composition of the risky portfolio.

9.3 Reduction to the wealth-tax-only case

Corollary 3 (Equivalence under the Norwegian system). *Under (C1)–(C3), the optimal portfolio under the combined tax system is identical to the optimal portfolio under a wealth tax alone. The portfolio distortion from non-uniform assessment is exactly*

$$\Delta \mathbf{w}^* = -\frac{\tau_w}{\gamma} \mathbf{V}^{-1}(\boldsymbol{\alpha} - \alpha_0 \mathbf{1})$$

(Frøseth 2026a, Proposition 5), regardless of the flow-tax rates τ_c and τ_d .

In the Norwegian system, where $\tau_k = \tau_c = 22\%$ and $r_s \approx r_f$, the flow taxes contribute zero portfolio distortion. The entire distortion comes from non-uniform wealth tax assessment—the result of Frøseth (2026a) applies without modification.

10 The Norwegian Tax System

The general framework of Sections 3 to 5 applies to any tax system with corporate, capital income, dividend, and wealth tax components. We now calibrate the framework to the Norwegian dual income tax (*aksjonærmodellen*), evaluating each of conditions (C1)–(C3) against the institutional design and quantifying distortion magnitudes under current parameters.

10.1 The dual income tax and conditions (C1)–(C2)

The Norwegian tax system taxes capital income at a flat rate of $\tau_k = 22\%$, identical to the corporate tax rate $\tau_c = 22\%$ (NOU 2014:13, 2014; Christiansen, 2004). Condition (C1) is therefore satisfied exactly—not approximately or as a policy aspiration, but as a structural feature of the dual income tax design. The Scheel Committee (NOU 2014:13, 2014) established this symmetry, and the Torvik Committee (NOU 2022:20, 2022) reaffirmed it as a cornerstone of the Norwegian system.

For the elevated tax on shareholder income above the normal return, the Norwegian system uses an *upward adjustment factor* (*oppjusteringsfaktor*) $f = 1.72$. The effective dividend and capital gains tax rate is

$$\tau_d = f \times \tau_k = 1.72 \times 0.22 = 0.3784. \quad (38)$$

This mechanism is the Norwegian implementation of the two-tier personal tax structure: rather than imposing a separate rate on excess returns, the system multiplies the income base by f and taxes at the ordinary rate τ_k . The factor f does not appear in conditions (C1)–(C3) and does not affect portfolio neutrality; it determines only the magnitude of the drift rescaling through the pass-through factor

$$k = (1 - \tau_c)(1 - \tau_d) = (1 - 0.22)(1 - 0.3784) = 0.485. \quad (39)$$

The gross-earnings multiplier is $1/k \approx 2.06$: for every krone of wealth tax paid from corporate earnings, the firm must generate 2.06 kroner in pre-tax profits.

The shielding deduction (*skjermingsfradrag*) sets r_s on the basis of the average three-month treasury bill rate plus a markup of 0.5 percentage points (see Section 7 for details). For 2025, the pre-tax basis is $r_s = 4.6\%$, corresponding to a published after-tax *skjermingsrente* of $r_s(1 - \tau_k) = 3.6\%$.³ Because the markup ensures $r_s \geq r_f$ whenever the treasury bill rate approximates the risk-free rate, condition (C2) is satisfied or slightly over-satisfied. The residual distortion from any discrepancy is bounded by

$$|(\tau_d - \tau_k)(r_s - r_f)| = |0.3784 - 0.22| \times |r_s - r_f| \approx 0.16 \times |r_s - r_f|.$$

With the 2025 markup of 50 basis points, this is 0.08%—negligible compared to the distortion from non-uniform assessment.

10.2 Valuation discounts and condition (C3)

Condition (C3)—uniform wealth tax assessment—is violated. The Norwegian wealth tax applies asset-class-specific valuation discounts, summarised in Table 5.

Table 5: Wealth tax assessment fractions in Norway (2025–2026).

Asset class	Base	Discount	Effective α_i
Bank deposits	Market	0%	1.00
Secondary housing	Assessed	0%	≈ 1.00
Holiday homes (<i>fritidsbolig</i>)	Assessed	70%	0.30
Listed shares	Market	20%	0.80
Unlisted shares	Book	20%	$0.80 \times B_i/M_i$
Commercial real estate	Assessed	20%	≤ 0.80
Primary housing ≤ 10 M NOK	Assessed	75%	0.25
Primary housing > 10 M NOK (excess)	Assessed	30%	0.70

The assessment fractions span a range of $\alpha_{\max} - \alpha_{\min} = 0.75$, from bank deposits ($\alpha = 1.00$)

³Source: Skatteetaten, *Skjermingsrente for aksjer og enkeltpersonforetak*, inntektsåret 2025.

to primary housing below the threshold ($\alpha = 0.25$). Listed shares receive a 20% discount on market value ($\alpha = 0.80$), placing them between the two extremes. Holiday homes (*fritidsbolig*) are assessed at 30% of estimated market value ($\alpha = 0.30$), creating an incentive comparable to primary housing.

The discount schedule has varied substantially over time. Shares carried no statutory valuation discount at all before 2017, received a 10% discount in 2017–2018, a 45% discount in 2021, a 25% discount in 2022, and the current 20% from 2023 onward (NOU 2022:20 2022, Table 10.3). Each change has a first-order effect on the (C3) distortion: the tightening from 45% to 20% roughly tripled the effective wealth tax on listed equity. Housing valuations are based on a hedonic regression model introduced in 2010, which uses transaction data to estimate market values annually. The 75% discount on primary housing has remained stable throughout and reflects a political compromise that preserves market-based assessment while shielding homeowners from the full tax impact.

The assessment system also applies proportional debt reduction (*gjeldsreduksjon*): debts secured against a given asset class are reduced by the same fraction $\beta_i = \alpha_i$ as the asset itself. This prevents taxpayers from combining the full debt deduction with discounted asset values, but it also means the effective discount extends to the net-of-debt position, reinforcing the tilt toward low- α assets (Frøseth 2026a, §4.3).

A critical distinction arises for *unlisted shares*. These are assessed at 80% of the company’s tax-assessed net asset value (*skattemessig formuesverdi*)—essentially book value—rather than market value. The effective assessment fraction is therefore $\alpha_i = 0.80 \times B_i/M_i$, where B_i is book value and M_i is market value. For firms with substantial intangible assets, goodwill, or growth options, the ratio B_i/M_i can be well below unity, making the effective assessment fraction much lower than 0.80. This is the book-value taxation channel identified in Frøseth (2026b), §9.2: the wealth tax base becomes a deterministic liability rather than a proportional claim on the stochastic market value, breaking the uniformity of the drift shift.

The resulting portfolio distortion follows from the formula of Frøseth (2026a), Proposition 5:

$$\Delta \mathbf{w}^* = -\frac{\tau_w}{\gamma} \mathbf{V}^{-1}(\boldsymbol{\alpha} - \alpha_0 \mathbf{1}).$$

Since $\alpha_{\text{listed}} < \alpha_{\text{deposits}}$, the wealth tax tilts portfolios toward shares and away from deposits. Since α_{unlisted} can be well below α_{listed} (depending on the book-to-market ratio), the tilt toward unlisted shares can be substantially larger than toward listed shares. Since $\alpha_{\text{housing}} < \alpha_{\text{listed}}$, there is a further tilt toward primary housing. The combined effect is an asset allocation that favours owner-occupied housing and unlisted equity in asset-light firms over bank deposits and listed shares.

10.3 Distortion magnitudes under current parameters

Table 6 compares the distortion magnitudes from each condition.

For a risk aversion parameter $\gamma = 3$, the (C3) distortion is of order 0.25% per unit of inverse

Table 6: Distortion magnitudes under the Norwegian tax system.

Condition	Parameter values	Distortion bound	Order
(C1): $\tau_k \neq \tau_c$	$\tau_k = \tau_c = 22\%$	$ \tau_k - \tau_c \cdot r_f = 0$	Exactly zero
(C2): $r_s \neq r_f$	$ r_s - r_f \leq 50$ bp	$0.16 \times 0.005 = 0.08\%$	Negligible
(C3): $\alpha_i \neq \alpha_j$	$\Delta\alpha \leq 0.75$	$\tau_w \cdot \Delta\alpha/\gamma$	$0.75\%/\gamma$

covariance—roughly 300 times larger than the upper bound on the (C2) channel, and the only nonzero channel among the three.

To put this in portfolio terms, consider the two-asset case of listed shares versus bank deposits. With $\tau_w = 1.0\%$, $\alpha_{\text{listed}} = 0.80$, $\alpha_{\text{deposits}} = 1.00$, and $\gamma = 3$:

$$\Delta w^* = \frac{0.01 \times 0.20}{3 \times \sigma^2} = \frac{0.00067}{\sigma^2}.$$

For an annual volatility of $\sigma = 0.20$ (typical for the Oslo Stock Exchange), this gives $\Delta w^* \approx 1.7$ percentage points toward listed shares.

For unlisted shares with a book-to-market ratio of $B/M = 0.5$ (representative of asset-light or high-growth firms), the effective assessment is $\alpha_{\text{unlisted}} = 0.80 \times 0.5 = 0.40$, and the tilt versus deposits is roughly four times larger: $\Delta w^* \approx 5$ percentage points. This illustrates why the book-value assessment of unlisted equity, identified in Frøseth (2026b), §9.2, creates a stronger portfolio incentive than the statutory 20% discount alone suggests.

For primary housing versus deposits, with $\Delta\alpha = 0.75$:

$$\Delta w^* = \frac{0.01 \times 0.75}{3 \times \sigma_h^2},$$

which for a housing volatility of $\sigma_h = 0.10$ yields $\Delta w^* \approx 25$ percentage points. The valuation discount on primary housing remains the dominant source of portfolio distortion in the Norwegian system (Frøseth 2026a, §4.3), but the book-value channel for unlisted shares is a significant secondary source whose magnitude depends on the firm’s asset composition.

10.4 Beyond the proportional framework

Several features of the Norwegian system fall outside the proportional framework of Section 4.

Shielding accumulation. Unused skjermingsfradrag carries forward and compounds at the shielding rate r_s . An investor holding shares with cost basis B for n years without fully utilising the deduction accumulates approximately $B[(1+r_s)^n - 1]$ in unused shielding. This accumulated shielding is *lost upon realisation*: it cannot be transferred to other shares or set against other income. The loss-upon-realisation creates a lock-in effect that makes the effective tax rate path-dependent and is not captured by the static framework (Auerbach, 1991). Quantifying this channel requires a dynamic model with heterogeneous holding periods.

Progressive wealth tax. The Norwegian wealth tax has a threshold of 1.9M NOK (sin-

gle, 2026) and a two-bracket rate structure: 1.0% on wealth between 1.9 M and 21.5 M NOK, and 1.1% above 21.5 M NOK. This introduces a wealth-dependent effective rate $\tau_w(W)$, creating a state-dependent drift modification that breaks the uniformity of the drift shift. In the Fokker–Planck framework, this is exactly the confining potential analysed in Frøseth (2026c): the progressive structure generates a restoring force that compresses the wealth distribution. The threshold also creates a tax shield in the sense of Frøseth (2026a), Proposition 6, which increases risk-taking for investors near the exemption boundary.

Retention and deferral. If a firm retains earnings, the dividend tax is deferred until distribution. This creates a timing option whose value depends on payout policy, future tax rates, and the investor’s discount rate. The effective tax rate becomes path-dependent, introducing a channel that favours retention and complicates the uniform-rescaling result. The participation exemption amplifies this deferral channel.

The participation exemption (*fritaksmetoden*). Corporate shareholders—aksjeselskaper, foreninger, and stiftelser—are exempt from tax on dividends and capital gains under the *fritaksmetoden*, introduced as part of the 2006 tax reform to prevent chain taxation (*kjedebeskatning*) of corporate profits through ownership layers (NOU 2022:20 2022). A partial correction was added in 2009: 3% of exempt income is included in ordinary income, yielding an effective tax rate of $0.03 \times 0.22 = 0.66\%$ on inter-company distributions. Dividends within a tax group (ownership above 90%) are fully exempt. Since most substantial Norwegian investors hold equities through unlisted holding companies, the *fritaksmetoden* is the dominant ownership structure for taxable wealth. A consequence, discussed in Frøseth (2026b), §9.2, is that the wealth tax base for these investors is assessed at beginning-of-period book value (1 January of the income year) rather than end-of-period market value—a predetermined quantity that simplifies the pricing equation and is empirically important for calibration.

For the framework of this paper, the *fritaksmetoden* changes the effective pass-through factor. A personal investor holding shares directly faces $k = (1 - \tau_c)(1 - \tau_d) \approx 0.485$. The same investor routing ownership through a holding company faces $k_{\text{hold}} \approx (1 - \tau_c)(1 - 0.0066) \approx 0.775$ on reinvested corporate earnings, because the dividend tax is replaced by the 0.66% effective rate within the corporate sector. The personal dividend tax is deferred until the investor extracts funds from the holding company. The result is a two-tier system: the tangency portfolio is preserved (the rescaling is still uniform across assets), but the *magnitude* of the rescaling differs between direct and indirect ownership, creating an incentive to interpose holding structures. Bjerksund and Schjelderup (2021a) analyse the neutrality of this two-tier system and show that the investor is indifferent between direct and indirect ownership only when the shielding rate equals the investor’s borrowing rate—a condition that generally does not hold, since the institutional r_s is based on the three-month treasury bill rate plus a 0.5 percentage point markup (see Section 7), which lies below typical borrowing rates. Bjerksund and Schjelderup (2021b) document the resulting lock-in: retained capital in Norwegian holding structures grew from approximately 500 M to 2,600 M NOK between 1999 and 2019, and they use a Black–Scholes framework to value the embedded option in the shielding system, finding forward equity tax credits of 0.36–1.34 NOK per krone invested depending on holding period and firm maturity.

The combination of the participation exemption and the shareholder model favours established investors with holding structures—who can defer personal dividend tax indefinitely—over those holding shares directly.

The quantitative significance is documented by [Bjerk Sund et al. \(2024\)](#), who compute effective tax rates for a tax-minimising investor holding unlisted shares through a holding structure. Using the average effective corporate tax rate of 13.1% (2004–2018) and an effective valuation discount of 65% on unlisted shares—reflecting the book-value assessment channel of Section 10.2, not just the statutory *verdsettelsesrabatt*—they find an overall effective tax rate of approximately 14.4%. This is well below the 21.4% for a passive investor in listed shares and the 27.3% average for a single worker. The *fritaksmetoden* accounts for the bulk of the gap: by eliminating the personal-level dividend tax on retained earnings, it reduces the effective burden from the combined tax system by roughly one-third for investors who can defer distributions indefinitely.

Wealth tax deferral. From 2026, Norway introduces a deferral scheme (*utsettelsesordning*) allowing taxpayers with taxable wealth above the threshold to defer wealth tax payments for up to three years, with interest accruing at the market rate. In the framework of Section 6, this converts the immediate liquidity drain into a deferred obligation, reducing the payment-channel amplification for constrained investors. The deferral does not change the present-value tax burden—interest ensures NPV-equivalence—but it alters the timing of the forced asset liquidation, which may affect portfolio dynamics in practice.

11 Discussion

11.1 Connection to the paper series

The results of this paper extend the neutrality framework of Papers 1–4 in a natural way. Table 7 summarises the connections.

Table 7: Connections to the paper series.

Prior result	Extension here
Paper 1, Prop. 2: Portfolio invariance (GBM)	Extends to combined flow + stock taxes
Paper 1, Prop. 3: Orthogonality	Combined taxes = uniform rescaling + shift
Paper 2, Prop. 1: Stochastic vol. neutrality	Same mechanism: CRRA absorbs flow taxes
Paper 2, Prop. 5: Non-uniform assessment	Remains the binding constraint
Paper 3, Def. 1: Drift-shift transformation	Generalises to drift-shift-and-rescale
Paper 3, Prop. 2: Neutrality as invariance	Relative drifts scaled by $(1 - \tau_c)(1 - \tau_d)$
Paper 4, Prop. 1: Gini preservation	Extends: uniform rescaling preserves Gini
Paper 7, Thm. 1: Spectral invariance	Pass-through factor k is isotropic perturbation

The connection to Paper 7 deserves elaboration. [Frøseth \(2026e\)](#) develops spectral portfolio theory by identifying neural network weight matrices as portfolio allocation matrices and proving that any *isotropic* perturbation to the portfolio objective preserves the singular-value distribution up to scale and shift (Theorem 1). The pass-through factor $k = (1 - \tau_c)(1 - \tau_d)$ of

the present paper is exactly such an isotropic perturbation: it rescales all excess returns uniformly, treating every asset symmetrically. The drift-shift-and-rescale symmetry of Theorem 1 is therefore the scalar projection of Paper 7’s spectral invariance. Conversely, the symmetry-breaking classification of Section 5 maps onto Paper 7’s isotropic/anisotropic distinction: conditions (C1) and (C2) ensure that the flow-tax rescaling is isotropic, while condition (C3) is the non-anisotropy requirement for the wealth tax—the same condition that prevents spectral distortion in the matrix-valued framework.

11.2 Equilibrium prices under inelastic markets

Frøseth (2026a), §5 analyses how the wealth tax affects equilibrium prices when asset supply is inelastic, following the inelastic markets hypothesis of Gabaix and Koijen (2022). The conclusion is that a uniform wealth tax depresses the price level of risky assets but does not change relative prices, because the drift shift is the same for all assets.

The same logic extends to the combined tax system. Under (C1)–(C3), the flow taxes rescale all excess returns by the uniform factor $k = (1 - \tau_c)(1 - \tau_d)$. This reduces the after-tax excess return on every risky asset by the same proportion, lowering aggregate demand for risky assets. Under inelastic supply, the equilibrium price level falls. But because the rescaling is uniform, relative prices are unchanged—the portfolio composition is unaffected.

The combined effect on the price level is larger than under the wealth tax alone: the wealth tax shifts the drift by $-\tau_w\alpha$, and the flow taxes compress the remaining excess drift by the factor k . Both effects reduce demand, and they compound. However, neither modifies relative prices, so pricing neutrality in the cross-sectional sense is preserved. This extends the result of Frøseth (2026a), §5 to the combined tax system.

11.3 The role of the shielding deduction

The shielding deduction has received attention in the tax policy literature primarily as an equity device: it ensures that a “normal” return on capital is not taxed at the elevated shareholder rate (Sørensen 2005). The present analysis adds a complementary perspective from allocative efficiency.

In the Fokker–Planck framework, the shielding deduction is the mechanism that aligns the boundary between the two personal tax regimes with the risk-free rate, converting a potentially distortionary two-regime tax structure into a uniform rescaling of excess returns. Without it, the dividend tax would apply to the *entire* after-corporate-tax return, creating an asymmetry between equity and debt that would distort the equity–debt allocation (though not the tangency portfolio among equities). With it, the only remaining source of portfolio distortion is non-uniform wealth tax assessment.

This provides a formal efficiency argument for the shielding deduction that is independent of the equity argument.

11.4 Policy implications

The Norwegian calibration of Section 10 yields a clear policy ranking: equalising assessment fractions across asset classes would do far more for allocative neutrality than any reform of the flow-tax rates or the shielding mechanism. The (C3) channel generates distortions roughly 300 times larger than the residual (C2) channel and is the only nonzero source of portfolio distortion in the Norwegian system. This echoes the empirical findings of [Jakobsen et al. \(2020\)](#) and [Brühlhart et al. \(2022\)](#), who document substantial portfolio responses to wealth tax incentives in Scandinavian and Swiss data.

11.5 Limitations

Beyond the Norwegian-specific qualifications discussed in Section 10.4 (shielding accumulation, progressive brackets, retention and deferral), one general limitation deserves mention.

Discrete time. In discrete time, the scaling factor $(1 - \tau_c)(1 - \tau_d)$ may affect the total risky allocation under non-CRRA preferences, even though it preserves the tangency portfolio. The magnitude of this effect for realistic parameters is an empirical question.

12 Conclusion

The combined system of ownership taxes—corporate tax, capital income tax, dividend tax, and wealth tax—preserves portfolio neutrality under three conditions: equal rates on capital income and corporate profit, alignment of the shielding rate with the risk-free rate, and uniform wealth tax assessment.

The drift-shift symmetry of [Frøseth \(2026d\)](#) generalises to a drift-shift-and-rescale symmetry: the flow taxes uniformly rescale the excess drift velocities by $(1 - \tau_c)(1 - \tau_d)$, while the wealth tax uniformly shifts all drifts. Neither modification alters relative drifts between assets. The distortions from flow taxes and from the wealth tax are additively separable and do not interact.

The shielding deduction emerges as a symmetry-restoring mechanism in the Fokker–Planck framework. By aligning the boundary between the two personal tax regimes with the risk-free rate, it converts a potentially distortionary two-regime structure into a uniform rescaling of excess returns. This provides a formal efficiency argument for the shielding deduction that complements the standard equity justification.

Calibrated to the Norwegian dual income tax, the framework confirms that flow-tax neutrality holds by design: $\tau_k = \tau_c = 22\%$ and the shielding rate tracks the risk-free rate. The pass-through factor $k = (1 - \tau_c)(1 - \tau_d) \approx 0.485$ compresses excess drifts but preserves their ranking. The dominant—and effectively the only—source of portfolio distortion is non-uniform wealth tax assessment, with valuation discounts ranging from 0% (bank deposits) to 75% (primary housing). The resulting portfolio tilts are roughly 300 times larger than the residual distortion from any shielding-rate misalignment. Equalising assessment fractions would do more for allocative neutrality than any reform of the flow-tax structure.

Acknowledgements

The author acknowledges the use of Claude (Anthropic) for assistance with literature review, L^AT_EX 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

- Alan J. Auerbach. Retrospective capital gains taxation. *American Economic Review*, 81(1): 167–178, 1991.
- Petter Bjerksund and Guttorm Schjelderup. Er den norske aksjonærbeskattningen nøytral? *Samfunnsøkonomen*, 135(4):43–52, 2021a.
- Petter Bjerksund and Guttorm Schjelderup. Aksjonærmodellen og fritaksmetoden: Et to-hodet troll? *Samfunnsøkonomen*, 135(4):53–72, 2021b.
- 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.
- Petter Bjerksund, Arnt Ove Hopland, and Guttorm Schjelderup. The taxation of Norway's richest. *Nordic Tax Journal*, 2024(s1):72–82, 2024. doi: 10.2478/ntaxj-2024-0006.
- Robin Boadway and Neil Bruce. A general proposition on the design of a neutral business tax. *Journal of Public Economics*, 24(2):231–239, 1984.
- Jean-Philippe Bouchaud and Marc Mézard. Wealth condensation in a simple model of economy. *Physica A*, 282(3–4):536–545, 2000.
- Marius Brühlhart, Jonathan Gruber, Matthias Krapf, and Kurt Schmidheiny. Behavioral responses to wealth taxes: Evidence from Switzerland. *American Economic Journal: Economic Policy*, 14(4):111–150, 2022. doi: 10.1257/pol.20200258.
- Vidar Christiansen. Norwegian income tax reforms. *ifo DICE Report*, 2(3):9–14, 2004.
- Evsey D. Domar and Richard A. Musgrave. Proportional income taxation and risk-taking. *Quarterly Journal of Economics*, 58(3):388–422, 1944.
- Anders G. Frøseth. Extensions to the wealth tax neutrality framework. 2026a. arXiv:[2603.05277](https://arxiv.org/abs/2603.05277) [physics.soc-ph].
- Anders G. Frøseth. Asset returns, portfolio choice, and proportional wealth taxation. 2026b. arXiv:[2603.05264](https://arxiv.org/abs/2603.05264) [physics.soc-ph].
- Anders G. Frøseth. From gravity to confinement: Wealth redistribution as optimal drift design in the Fokker–Planck framework. Working paper, 2026c.
- Anders G. Frøseth. Wealth taxation as a drift modification: A Fokker–Planck approach to tax neutrality. 2026d. arXiv:[2603.05283](https://arxiv.org/abs/2603.05283) [physics.soc-ph].

- Anders G. Frøseth. Spectral portfolio theory: From SGD weight matrices to wealth dynamics. 2026e. arXiv:[2603.09006](https://arxiv.org/abs/2603.09006) [physics.soc-ph].
- Xavier Gabaix and Ralph S. J. Koijen. In search of the origins of financial fluctuations: The inelastic markets hypothesis. Working paper, Harvard University and University of Chicago, 2022.
- Katrine Jakobsen, Kristian Jakobsen, Henrik Kleven, and Gabriel Zucman. Wealth taxation and wealth accumulation: Theory and evidence from Denmark. *Quarterly Journal of Economics*, 135(1):329–388, 2020.
- James Mirrlees, Stuart Adam, Timothy Besley, Richard Blundell, Stephen Bond, Robert Chote, Malcolm Gammie, Paul Johnson, Gareth Myles, and James Poterba. *Tax by Design: The Mirrlees Review*. Institute for Fiscal Studies, London, 2011.
- NOU 2003:9. Skatteutvalget: Forslag til endringer i skattesystemet. Technical report, Norges offentlige utredninger, Oslo, 2003.
- NOU 2014:13. Kapitalbeskatning i en internasjonal økonomi. Technical report, Norges offentlige utredninger, Oslo, 2014.
- NOU 2022:20. Et helhetlig skattesystem. Technical report, Norges offentlige utredninger, Oslo, 2022.
- Agnar Sandmo. Portfolio theory, asset demand and taxation: Comparative statics with many assets. *Review of Economic Studies*, 44(2):369–379, 1977.
- Peter Birch Sørensen. Neutral taxation of shareholder income. *International Tax and Public Finance*, 12(6):777–801, 2005.
- Joseph E. Stiglitz. The effects of income, wealth, and capital gains taxation on risk-taking. *Quarterly Journal of Economics*, 83(2):263–283, 1969.