Combined Flow and Stock Taxation under Heterogeneous Returns: A Cross-sectional and Long-run Analysis of the Norwegian System

Anders G. Frøseth*

July 8, 2026

Abstract

Under three structural conditions on the joint flow and stock tax system a personal shareholder simultaneously faces, the post-tax mean-variance optimum is shifted by the same direction in asset space, leaving the cross-sectional dispersion of portfolios spectrally invariant (Frøseth, 2026b,h). The actual Norwegian system violates the underlying assumptions on four fronts: heterogeneous expected returns (Fagereng et al., 2020), non-uniform wealth-tax assessment (Frøseth, 2026a), leveraged household balance sheets, and a progressive bracket schedule. We lift all four on a single deterministic-equilibrium framework.

Theorem 1 gives the closed-form post-tax shift on the cross-sectional allocation matrix under heterogeneous returns: a multiplicative half compressing cross-agent dispersion plus a rank-one additive half along an assessment-weighted direction, with rank scaling $\min(d+1, N)$ under d-factor ability stratification. *Theorem 2* gives the closed-form type-conditional Pareto exponent of the long-run wealth distribution under heterogeneous returns and differential assessment.

Six corollaries lift the simplifying assumptions to the post-2017 Norwegian dual income tax: the empirical assessment profile (Cor. 4); canonical leverage with *gjeldsreduksjon* and corporate routing under *fritaksmetoden* (Cor. 5: the leverage matrix Λ enters the FOC via Λ-conjugation, and the per-equity wealth-tax wedge is unchanged); K-bracket progressive schedules adding K rank-one bracket-indicator modes to the cross-sectional spectrum (Cor. 6). A proposition extends Theorem 2 with a policy-invariance reading: the upper-tail Pareto exponent depends only on the topmost bracket rate.

Three comparative-statics findings emerge: the Pareto exponent is non-monotonic in the assessment wedge and peaks at the agent's-portfolio-zero point; Norway's actual profile sits near this peak; and the cross-type Pareto spread changes sign with ability, with anisotropic assessment compressing the bottom of the wealth distribution and concentrating the top. A unified Norwegian worked example illustrates all four channels.

JEL Classification: G11, G12, H21, H22, H24, C65.

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

Keywords: combined flow and stock taxes, heterogeneous returns, ability heterogeneity, differential wealth-tax assessment, leveraged household balance sheet, progressive wealth taxation, allocation matrix, rank scaling, type-conditional Pareto exponent, joint Fokker–Planck, Norwegian dual income tax.

Contents

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, leaving the investor's optimal risk exposure 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. Frøseth (2026e) sharpens the wealth-tax case, showing that a uniform levy on the market value of all assets leaves portfolio weights, Sharpe ratios, and asset prices unchanged; Frøseth (2026g) recasts the result as a *drift-shift symmetry* in the Fokker–Planck equation for wealth dynamics, with the proportional wealth tax acting as a uniform translation of the drift coefficient that leaves the diffusion structure of the wealth process entirely intact.

A real shareholder, however, does not face the wealth tax in isolation. A Norwegian taxpayer collecting dividends from her corporate equity holdings is subject to four distinct levies, collectively the *eierkostnader* (ownership costs): a *corporate tax* at rate τ_c on gross profits; a *capital-income tax* at rate τ_k on the risk-free return; a *dividend and capital gains tax* at rate τ_d on the excess return; and a *wealth tax* at rate τ_w on net assets. Frøseth (2026b) extends the neutrality framework to this combined system and shows that portfolio neutrality is preserved under three structural conditions:

(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 fraction is uniform across assets: $\alpha_i = a$ for all i.

Under (C1)–(C3) the after-tax excess return on every risky asset is a uniform rescaling of its pre-tax counterpart by a single factor λ, together with a row-uniform shift c, while the diffusion structure of the wealth process is preserved. The drift-shift symmetry of the wealth-tax-only case generalises to a *drift-shift-and-rescale symmetry* on the joint flow and stock tax system (Frøseth, 2026b).

Frøseth (2026h) establishes a cross-sectional companion to this neutrality. Stack the optimal portfolio weights of T households across N risky assets into an *allocation matrix* $\mathbf{W} \in \mathbb{R}^{T\times N}$, with row t the weight vector of household t and column j the cross-household distribution of holdings of asset j. The singular-value distribution of $\mathbf{W}$ —its *spectral structure*—encodes how diversified the cross-section of households is, what cross-sectional factor portfolios dominate, and how concentrated holdings are across investors. Under (C1)–(C3) and the simplifying assumption that all households share a common expected-return vector $\boldsymbol{\mu} \in \mathbb{R}^N$, every household's mean-variance optimum is shifted post-tax by the same vector $\boldsymbol{\delta} \in \mathbb{R}^N$ in asset space. The cross-row dispersion of $\mathbf{W}$ is therefore preserved up to a row-uniform translation, and the singular-value distribution is invariant. Frøseth (2026h) calls this the *cross-sectional neutrality* or *spectral neutrality* corollary of (C1)–(C3).

The assumption that all households share a common expected-return vector is empirically tenuous. Fagereng et al. (2020), using Norwegian household-portfolio register data, document persistent dispersion in returns across the wealth distribution: investors at the top of the wealth distribution earn excess returns of order 200 basis points relative to the median, with the gap stable across decades and consistent with skill, scale, and information-access differences rather than *ex post* luck. A recent corpus literature absorbs this empirical fact and develops its policy consequences. Frøseth (2026c) introduces an ability-stratified extension of the Fokker–Planck framework: each household carries an individual ability parameter z that scales its drift vector $\boldsymbol{\mu}(z)$, so that a uniform drift reduction disproportionately burdens low-ability investors and breaks the drift-shift neutrality of the homogeneous-return baseline. Frøseth (2026d) pushes the line further, deriving a closed-form joint Pareto exponent for the (log-wealth, ability) distribution under the full (C1)–(C3) flow and stock tax system, with explicit dependence on flow-tax, wealth-tax, and demographic-turnover channels. The cross-sectional spectral question—does the (C1)–(C3) system still preserve the singular-value distribution of the allocation matrix when households face heterogeneous returns?—has not been settled.

This paper closes the gap between the homogeneous-returns spectral neutrality result of Frøseth (2026h) and the actual Norwegian dual income tax, lifting four channels on a single deterministic-equilibrium framework.

Channel 1: heterogeneous returns. Under (C1)–(C3) and heterogeneous returns the post-tax shift in the allocation matrix admits a closed-form additive decomposition into a *row-by-row rescaling* of each household's excess-return signal $\boldsymbol{\mu}_t - r_f\mathbf{1}$ by the factor $\lambda - 1$ and a *row-uniform translation* toward the global-minimum-variance portfolio direction (Theorem 1). When returns admit an ability-affine decomposition $\boldsymbol{\mu}_t = \boldsymbol{\mu}_0 + \sum_{k=1}^{d} z_{t,k}\,\boldsymbol{\nu}_k$ in d ability factors, the post-tax shift has rank exactly $d+1$ in the typical financial regime: one row-uniform direction recovered from Frøseth (2026h) plus one ability-driven distortion direction per ability dimension (Cor. 2). Spectral neutrality breaks at a rate equal to the effective dimension of ability heterogeneity. Setting $\boldsymbol{\mu}_t \equiv \boldsymbol{\mu}$ recovers Frøseth (2026h); setting $\tau_w = 0$ generates the heterogeneity-compression of Frøseth (2026c); setting $\tau_c = \tau_d = 0$ recovers the drift-shift symmetry of Frøseth (2026g) (Cor. 3).

Channel 2: differential assessment. The Norwegian wealth tax does not satisfy (C3): primary housing is assessed at $\alpha = 0.25$ of market value, listed shares and commercial property at $\alpha = 0.80$, and bank deposits at $\alpha = 1.00$ (Frøseth, 2026a, Table 2). Cor. 4 lifts (C3) and shows that the rank decomposition survives verbatim under differential assessment, with the additive direction shifting from $\mathbf{V}^{-1}\mathbf{1}$ to the assessment-weighted direction $\mathbf{V}^{-1}\mathbf{c}$ and the fluctuation-SDE invariance carrying through. The broken-(C3) coupling between the agent's portfolio and the assessment vector closes a long-run channel that Section 5 turns into a closed-form type-conditional Pareto exponent (Theorem 2) by substituting the broken-(C3) optimal portfolio into the joint Fokker–Planck framework of Frøseth (2026d). Three substantive comparative-statics findings follow: the Pareto exponent is non-monotonic in the assessment wedge (peaking where the agent's optimal portfolio is at zero risky exposure); the actual Norwegian assessment profile sits near this peak; and the cross-type Pareto-exponent spread between broken-(C3) and (C3)-uniform regimes changes sign with ability, with anisotropic Norwegian assessment compressing the bottom of the wealth distribution and concentrating the top.

Channel 3: leveraged household balance sheets. Norwegian households leverage the discounted asset classes heavily through the residential mortgage market and through holding-company structures. Under the post-2017 institutional configuration — proportional debt-reduction (*gjeldsreduksjon*) on the wealth-tax base and corporate-routed asset holding with uncapped interest deduction via corporate-loss carry-forward — the leveraged tax stack collapses onto the unleveraged structure of Cor. 4 up to a single transformation by a leverage matrix Λ (Cor. 5). Leverage acts on the FOC entirely through the gearing factor $1/(1 - \ell_j)$ on the gross-return premium; the wealth-tax wedge per unit of equity is unchanged.

Channel 4: progressive bracket schedules. Norway's post-2022 wealth tax is a two-bracket piecewise-constant function of net taxable wealth (rates 1.0% and 1.1%, threshold near NOK 20 M). Cor. 6 extends the leveraged FOC shift to a K-bracket schedule; the Norwegian $K = 2$ case picks up one extra rank degree of freedom relative to Cor. 5, observable as a threshold-crossing mode in the cross-sectional spectrum. Proposition 1 extends Theorem 2 to the progressive setting with a policy-invariance reading: the upper-tail Pareto exponent depends only on the topmost bracket rate, so changing the lower-bracket rate or the threshold leaves the upper tail unchanged.

Norwegian unified illustration. Under the post-2017 *fritaksmetoden*, conditions (C1) and (C2) hold by institutional design. Section 6 runs all four channels on a single calibration: the cross-sectional spectrum at $d = 2$, the type-conditional Pareto exponent across an ability range with the sign-changing cross-type spread, and the rank progression from broken (C3) through canonical leverage to the post-2022 progressive schedule.

The paper is organised as follows. Section 2 fixes the heterogeneous-returns setup: the allocation matrix, the row-varying mean-variance objective, the (C1)–(C3) combined-tax transformation, and the homogeneous-returns baseline of Frøseth (2026h). Section 3 states and proves the heterogeneous-returns FOC shift theorem (Theorem 1), the first main result. Section 4 draws out six corollaries: the homogeneous-returns reduction (Corollary 1), the rank-scaling result (Corollary 2), the flow-vs-stock sub-system limits (Corollary 3), the lift to broken (C3) under differential wealth-tax assessment (Corollary 4), the lift to leveraged household balance sheets under the canonical Norwegian configuration (Corollary 5), and the lift to progressive bracket schedules (Corollary 6). Section 5 states and proves the type-conditional Pareto-exponent transformation under broken (C3) and heterogeneous returns (Theorem 2, the second main result), reads off three comparative-statics findings, and extends to progressive schedules (Proposition 1). Section 6 presents the unified Norwegian worked example, illustrating all four channels on a single calibration. Section 7 traces the contribution into the corpus's broader architecture; Section 8 discusses implications and open directions; Section A collects the proofs of the six corollaries.

2 Setup

This section assembles the four pieces the analysis rests on: the cross-sectional allocation matrix (§2.1), the per-row mean-variance optimum that pins down its deterministic equilibrium (§2.2), the combined flow and stock tax operator that maps pre-tax to post-tax expected returns (§2.3), and the homogeneous-returns baseline against which the new result is benchmarked (§2.4). The allocation matrix and the loss follow Frøseth (2026h); the tax operator follows Frøseth (2026b); the row-level ability stratification of returns follows Frøseth (2026c).

2.1 The heterogeneous-returns allocation matrix

The central object is the cross-sectional *allocation matrix*: rows index households, columns index risky assets, and entry $\mathbf{W}_{t,j}$ is the portfolio weight household t places on asset j. Its singular-value distribution encodes how diversified the cross-section of households is, what cross-sectional factor portfolios dominate, and how concentrated holdings are across investors. We work in the deterministic-equilibrium frame of Frøseth (2026h), generalised so that each household carries an ability label z_t that determines its own drift vector $\boldsymbol{\mu}(z_t)$—different rows of $\mathbf{W}$ now optimise against different expected-return vectors, in the spirit of the ability-stratified setup of Frøseth (2026c).

Definition 1 (Heterogeneous-returns allocation matrix)**.** Let $T, N \geq 1$ be integers and $\mathcal{Z}$ a measurable space of agent abilities. A *heterogeneous-returns allocation matrix* is a tuple $(\mathbf{W}, \{z_t\}_{t=1}^{T}, \boldsymbol{\mu}(\cdot), \mathbf{V}, r_f, \gamma)$ consisting of an allocation matrix $\mathbf{W} \in \mathbb{R}^{T\times N}$, ability labels $z_t \in \mathcal{Z}$, an ability-to-drift map $\boldsymbol{\mu} : \mathcal{Z} \to \mathbb{R}^N$, a positive-definite covariance $\mathbf{V} \in \mathbb{R}^{N\times N}$ shared across rows, a risk-free rate $r_f \in \mathbb{R}$, and a risk-aversion coefficient $\gamma > 0$. We write $\boldsymbol{\mu}_t := \boldsymbol{\mu}(z_t)$ for the drift vector seen by agent t and collect the row drifts in the *drift matrix* $\mathbf{M} \in \mathbb{R}^{T\times N}$ defined by $\mathbf{M}_{t\cdot} := \boldsymbol{\mu}_t^{\top}$.

The setup of Frøseth (2026h) corresponds to the special case $\boldsymbol{\mu}_t \equiv \boldsymbol{\mu}$, in which $\mathbf{M} = \mathbf{1}_T\boldsymbol{\mu}^{\top}$

has rank one. The row covariance $\mathbf{V}$ is shared across agents throughout: this is the diffusion-invariance assumption already in force in Frøseth (2026b) and Frøseth (2026c), and it is preserved by the C1–C3 flow and stock tax system of Frøseth (2026b) regardless of return heterogeneity.

2.2 Row-varying mean-variance loss and first-order condition

Each row of $\mathbf{W}$ is the agent's risky-asset weight vector. The deterministic equilibrium $\bar{\mathbf{W}}$ is the matrix that minimises the household-by-household mean-variance loss; the first-order condition (henceforth FOC) for that minimum—obtained by setting the gradient of the loss with respect to each row to zero—then determines the optimal portfolio weights row by row. In excess-return form the mean-variance loss factorises across rows, with a row-varying linear term and a row-uniform quadratic curvature:

$$\mathcal{L}(\mathbf{W}) \;=\; -\sum_{t=1}^{T} \mathbf{W}_{t\cdot}^{\top}(\boldsymbol{\mu}_t - r_f\mathbf{1}) \;+\; \tfrac{\gamma}{2}\sum_{t=1}^{T} \mathbf{W}_{t\cdot}^{\top}\mathbf{V}\,\mathbf{W}_{t\cdot}. \tag{1}$$

The deterministic FOC for row t (zero of the gradient $\partial\mathcal{L}/\partial\mathbf{W}_{t\cdot}$) is $\bar{\mathbf{W}}_{t\cdot} = (1/\gamma)\,\mathbf{V}^{-1}(\boldsymbol{\mu}_t - r_f\mathbf{1})$. Stacking rows,

$$\bar{\mathbf{W}} \;=\; \tfrac{1}{\gamma}\,(\mathbf{M} - r_f\mathbf{1}_T\mathbf{1}_N^{\top})\,\mathbf{V}^{-1}, \tag{2}$$

so $\operatorname{rank}(\bar{\mathbf{W}}) = \operatorname{rank}(\mathbf{M} - r_f\mathbf{1}_T\mathbf{1}_N^{\top})$. This rank takes any value in $\{1,\ldots,\min(T,N)\}$ depending on the dimension of the ability-induced heterogeneity in $\boldsymbol{\mu}(\cdot)$: one for the homogeneous baseline, two when $\boldsymbol{\mu}$ depends affinely on a scalar ability, and up to $\min(T,N)$ for fully unstructured row-by-row drift.

2.3 The combined-tax transformation $\mathcal{T}^{\text{gen}}$ on heterogeneous returns

The combined flow and stock tax system of Frøseth (2026b) acts on the risky-asset return process via four taxes — corporate τ_c, dividend τ_d, and capital-income τ_k on the flow, and a wealth tax τ_w on the stock — together with an assessment-fraction profile $\{\alpha_i\}_{i=0}^{N}$ for the wealth tax (with α_0 the fraction applied to the risk-free asset). The neutrality conditions C1–C3 of Frøseth (2026b) require $\tau_k = \tau_c$ (C1), shielding rate $\rho_s = r_f$ (C2), and $\alpha_i = a$ uniform across risky assets (C3). We work throughout under C1–C3 with $\alpha_0 = 0$, and define the derived scalars

$$\lambda \;:=\; (1-\tau_c)(1-\tau_d), \qquad c \;:=\; -(1-\tau_d)\tau_c\,r_f \;-\; \tau_w\,a. \tag{3}$$

The flow-tax retention factor λ matches P5's notation; in the corpus-canonical conventions of Frøseth (2026f,d) the same quantity is denoted k and we identify $\lambda \equiv k$ throughout. The scalar $\lambda \in (0,1]$ is the after-tax multiplicative compression of every excess return, encoding the flow-tax channel (corporate τ_c and dividend τ_d); c is a uniform additive drag, combining the shielding-corrected after-flow-tax discount on the risk-free anchor with the wealth-tax stock component. The four-tax book-keeping thus collapses to the pair (λ, c), and the heterogeneous-returns analysis below acts directly on these two parameters. Under these conditions Frøseth

(2026b, §4) shows that the after-tax excess return on risky asset j for agent t is

$$R_{t,j}^{\text{ex,post}} \;=\; \lambda\,(\mu_{t,j} - r_f) \;+\; c, \tag{4}$$

i.e. a row-pointwise affine transformation of the pre-tax excess return that is uniform across both assets and agents. The diffusion covariance $\mathbf{V}$ is preserved: this is the essential content of C3 combined with the asset-level diffusion-invariance argument of Frøseth (2026b), and it carries over verbatim under row heterogeneity because $\mathbf{V}$ is asset-level and ability-independent in the joint Fokker–Planck setup of Frøseth (2026c). We write $\mathcal{T}^{\text{gen}}$ for the combined-tax transformation acting on the underlying SDE so as to produce (4).

2.4 The homogeneous-returns baseline

Before turning to the heterogeneous-returns analysis, we record the homogeneous-returns baseline of Frøseth (2026h)—the limit case against which the new result is measured. When $\boldsymbol{\mu}_t \equiv \boldsymbol{\mu}$ for all t, the drift matrix $\mathbf{M} = \mathbf{1}_T\boldsymbol{\mu}^\top$ has rank one and (2) collapses to $\bar{\mathbf{W}}_{\text{hom}} = (1/\gamma)\,\mathbf{1}_T\,(\mathbf{V}^{-1}(\boldsymbol{\mu}-r_f\mathbf{1}))^\top$, itself rank one along the tangency direction $\mathbf{V}^{-1}(\boldsymbol{\mu}-r_f\mathbf{1})$. Substituting (4) into the row-FOC and subtracting the pre-tax FOC then yields the *homogeneous-returns FOC shift*

$$\bar{\mathbf{W}}^{\text{post}}_{\text{hom}} - \bar{\mathbf{W}}_{\text{hom}} \;=\; \mathbf{1}_T\,\boldsymbol{\delta}^\top, \qquad \boldsymbol{\delta} \;:=\; \tfrac{1}{\gamma}\,\mathbf{V}^{-1}\big((\lambda-1)(\boldsymbol{\mu}-r_f\mathbf{1}) + c\,\mathbf{1}\big), \tag{5}$$

a rank-one matrix factoring as $\mathbf{1}_T \otimes \boldsymbol{\delta}$. This is the C1–C3 tax-neutrality corollary of Frøseth (2026h): the shift acts on every row identically, so the cross-sectional spectral structure of $\mathbf{W}$ is preserved up to a row-uniform translation. The question this paper answers is what becomes of (5) when the row drifts $\boldsymbol{\mu}_t$ are allowed to vary, i.e. when $\mathbf{M}$ has rank above one.

2.5 Notation summary

Table 1 collects the symbols used throughout. The conventions follow the cross-paper notation reference of the corpus, with two paper-specific carry-overs from Frøseth (2026b,h): the flow-tax retention factor is denoted λ (equivalently k in Frøseth, 2026f,d) and the per-asset assessment is written α_i in component form, $\boldsymbol{\alpha}$ in vector form. Scalar (uniform) wealth-tax assessment uses Roman a to disambiguate from the type-conditional Pareto exponent $\alpha(z)$ of Theorem 2.

Table 1: Notation used in this paper. Type / ability variable z matches Frøseth (2026c); tax parameters $(\tau_c, \tau_d, \tau_w, r_f, r_s)$ follow the corpus-canonical naming.

Symbol	Meaning
Tax-system parameters	
τ_c, τ_d, τ_w	Corporate, dividend, wealth-tax rates
τ_k	Capital-income tax rate (set $\tau_k = \tau_c$ under (C1))
r_f, r_s	Risk-free rate, shielding rate ($r_s = r_f$ under (C2))
a	Wealth-tax assessment fraction (scalar; under (C3))
$\boldsymbol{\alpha}$, α_i	Per-asset assessment vector / component (broken (C3))
$\lambda = (1-\tau_c)(1-\tau_d)$	Flow-tax retention; equivalent to k in Frøseth, 2026d
c, $\mathbf{c}$	Additive shift (scalar under (C3); asset-vector under broken (C3))
Cross-sectional structures	
T, N, d	Households, risky assets, ability factors
$\mathbf{W}$, $\mathbf{M}$, $\mathbf{V}$	Allocation matrix, drift matrix, asset-level covariance
$\mathbf{M}_{\text{centered}} = \mathbf{M} - r_f \mathbf{1}_T \mathbf{1}_N^\top$	Risk-free-centred drift matrix
$\bar{\mathbf{W}}$, $\bar{\mathbf{W}}^{\text{post}}$	Pre- / post-tax FOC-equilibrium allocation matrix
γ	Risk-aversion coefficient
$\boldsymbol{\nu}_k$, $\mathbf{z}_k$	Ability-factor loadings / scores ($k = 1, \ldots, d$)
Leveraged balance sheet (Section 4.5)	
V_j, D_j, E_j	Gross asset value, debt, equity in class j; $E_j = V_j - D_j$
$W = \sum_j E_j$	Household net wealth (= total equity); static analogue of X_t
$\ell_j = D_j / V_j$	Loan-to-value (LTV) ratio on asset class j
$\Lambda = \operatorname{diag}(1 - \ell_j)$	Leverage matrix
$\theta_j = E_j / W$	Equity portfolio share for class j (used in Corollary 5)
$\mathbf{V}^{\text{eq}} = \Lambda^{-1} \mathbf{V} \Lambda^{-1}$	Equity-return covariance matrix
Progressive wealth-tax schedule (Section 4.6)	
K, $\{B_b\}_{b=1}^K$	Number of brackets and the brackets themselves
$\tau_w^{(b)}$	Wealth-tax rate in bracket b
$b(t)$, $\mathbf{1}^{(b)} \in \{0,1\}^T$	Bracket assignment of agent t and indicator vector
$\mathbf{c}^{(b)} \in \mathbb{R}^N$	Bracket-conditional shift vector (Equation (17))
Heterogeneous-returns / long-run primitives	
$z \in Z$	Type / ability variable (canonical, after Frøseth, 2026c)
$\boldsymbol{\mu}(z)$, $\sigma(z)$	Type-conditional return primitives
$m_0(z) = \mu(z) - \sigma(z)^2/2$	Type-conditional log-drift (canonical)
X_t, $Y_t = \log X_t$	Wealth, log-wealth processes
$w(z)$	Post-tax optimal portfolio fraction (single asset)
$m_X(z)$	Per-dollar wealth-process drift
$v_{\text{eff}}(z)$, $D_{\text{eff}}(z)$	Effective log-wealth drift / diffusion
$\delta(z)$	Type-conditional demographic turnover rate
$\alpha(z)$	Type-conditional Pareto exponent (*distinct from* a)
$\Delta\alpha(z)$	Cross-regime Pareto-exponent spread
Conventions	
$\mathbf{1}$, $\mathbf{1}_T$, $\mathbf{1}_N$	All-ones vectors
d	Differential operator d
$W_t^{(z)}$	Type-conditional standard Brownian motion (*distinct from* $\mathbf{W}$)

3 Heterogeneous-returns FOC shift

This section states and proves the central result of the paper. §3.1 states the post-tax shift in the deterministic-equilibrium allocation matrix as a sum of two contributions—a multiplicative-half rescaling each household's pre-tax excess-return signal, and an additive-half row-uniform translation along the global-minimum-variance direction—together with the rank bound $\min(T, N)+1$. §3.2 gives the proof, which follows by stacking the row-by-row mean-variance first-order condition and reading off the additive decomposition of the resulting drift-shift matrix. The two halves admit a clean tax-policy reading: the multiplicative half is the cross-sectional signature of the flow-tax stack's differential incidence, contracting cross-agent dispersion in optimal allocations proportional to $1-\lambda$; the additive half is the wealth-tax channel, translating every household uniformly along the global- minimum-variance direction.

3.1 Statement

Theorem 1 (Heterogeneous-returns FOC shift)**.** *Under the combined flow and stock tax system of Frøseth (2026b) satisfying (C1)-(C3) with $\alpha_0 = 0$, and under heterogeneous returns ($\boldsymbol{\mu}_t \in \mathbb{R}^N$ per row t, common $\mathbf{V}$), the deterministic equilibrium of the allocation matrix shifts by*

$$\bar{\mathbf{W}}^{\text{post}} - \bar{\mathbf{W}} = \frac{1}{\gamma}\,\Delta\mathbf{M}\,\mathbf{V}^{-1},$$

where $\Delta\mathbf{M} \in \mathbb{R}^{T\times N}$ has row t equal to

$$(\Delta\mathbf{M})_{t\cdot} = (\lambda-1)(\boldsymbol{\mu}_t - r_f\mathbf{1})^\top + c\,\mathbf{1}^\top.$$

The shift decomposes additively into:

- *A* multiplicative-half *contribution* $(\lambda-1)(\mathbf{M}-r_f\mathbf{1}_T\mathbf{1}_N^\top)\mathbf{V}^{-1}/\gamma$*, whose rank equals* $\operatorname{rank}(\mathbf{M}-r_f\mathbf{1}_T\mathbf{1}_N^\top) \in \{1,\dots,\min(T,N)\}$*.*
- *An* additive-half *contribution* $(c/\gamma)\,\mathbf{1}_T(\mathbf{V}^{-1}\mathbf{1})^\top$*, rank-one along the global-minimum-variance direction.*

Total rank: at most $\min(T,N)+1$*.*

Remark (Economic interpretation)*.* The two halves of Theorem 1 correspond to two economically distinct tax channels acting on the cross-section simultaneously. The multiplicative half scales each household's pre-tax tangency portfolio $\mathbf{V}^{-1}(\boldsymbol{\mu}_t - r_f\mathbf{1})/\gamma$ by $\lambda \in (0,1]$, contracting cross-agent dispersion in optimal allocations by factor $1-\lambda$: high-ability households—those with the largest pre-tax expected returns—bear the largest absolute flow-tax liability and respond by retreating proportionally toward the cross-sectional mean. The additive half is row- uniform: every household, regardless of ability, receives the same translation $(c/\gamma)\,\mathbf{V}^{-1}\mathbf{1}$ along the global-minimum-variance direction. This is the wealth-tax channel; it shifts the common location of the cross-section but leaves cross-agent dispersion unchanged. The rank bound $\min(T,N)+1$ caps the number of independent cross-sectional distortion channels: at most $\min(T,N)$ from the multiplicative half (one per linearly-independent column of pre-tax excess returns) and one more from the wealth-tax direction; Corollary 2 sharpens this to $\min(d+1, N)$ when returns admit a d-factor ability decomposition.

Remark (Absorption of the additive half). The bound $\min(T, N)+1$ is generally slack. Consider ability-affine returns

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_0 + \sum_{k=1}^{d} z_{t,k}\,\boldsymbol{\nu}_k, \qquad \mathbf{z}_k := (z_{1,k}, \ldots, z_{T,k})^\top \in \mathbb{R}^T,$$

covering the homogeneous baseline ($d = 0$), scalar ability ($d = 1$), and multi-factor ability lifts à la Frøseth (2026c,d). Column j of $\mathbf{M} - r_f \mathbf{1}_T \mathbf{1}_N^\top$ then equals $(\mu_{0,j} - r_f)\,\mathbf{1}_T + \sum_k \nu_{k,j}\,\mathbf{z}_k$, so the multiplicative-half's column span lies in $\operatorname{span}(\mathbf{1}_T, \mathbf{z}_1, \ldots, \mathbf{z}_d) \subset \mathbb{R}^T$; the additive-half's column span is $\operatorname{span}(\mathbf{1}_T)$. The $\mathbf{1}_T$ direction is therefore *shared* whenever $\boldsymbol{\mu}_0 \neq r_f \mathbf{1}$, since at least one column of the multiplicative-half then carries a non-zero $\mathbf{1}_T$ component. Under this typical financial parametrisation the additive-half is absorbed into the multiplicative-half's column span and the total FOC-shift rank equals the multiplicative-half rank, *not* multiplicative-half rank plus one. The $+1$ bump materialises only in the degenerate case $\boldsymbol{\mu}_0 = r_f \mathbf{1}$ (in particular the zero-baseline case $\boldsymbol{\mu}_0 = 0$ with $r_f = 0$, where the columns of $\mathbf{M}$ are pure ability combinations and exclude the $\mathbf{1}_T$ direction).

3.2 Proof of Theorem 1

Substitute the after-tax excess return formula (4) into the per-row mean-variance FOC (§2.2), replacing the pre-tax excess return $\boldsymbol{\mu}_t - r_f \mathbf{1}$ by its post-tax counterpart $\lambda(\boldsymbol{\mu}_t - r_f \mathbf{1}) + c\,\mathbf{1}$:

$$\bar{\mathbf{W}}^{\text{post}}_{t\cdot} \;=\; \tfrac{1}{\gamma}\,\mathbf{V}^{-1}\big(\lambda(\boldsymbol{\mu}_t - r_f \mathbf{1}) + c\,\mathbf{1}\big). \tag{6}$$

Subtracting the pre-tax row-FOC $\bar{\mathbf{W}}_{t\cdot} = (1/\gamma)\,\mathbf{V}^{-1}(\boldsymbol{\mu}_t - r_f \mathbf{1})$ gives the per-row shift

$$\bar{\mathbf{W}}^{\text{post}}_{t\cdot} - \bar{\mathbf{W}}_{t\cdot} \;=\; \tfrac{1}{\gamma}\,\mathbf{V}^{-1}\big((\lambda - 1)(\boldsymbol{\mu}_t - r_f \mathbf{1}) + c\,\mathbf{1}\big). \tag{7}$$

Stacking rows yields the matrix identity

$$\bar{\mathbf{W}}^{\text{post}} - \bar{\mathbf{W}} \;=\; \tfrac{1}{\gamma}\,\Delta\mathbf{M}\,\mathbf{V}^{-1},$$

where row t of $\Delta\mathbf{M} \in \mathbb{R}^{T \times N}$ equals $(\lambda-1)(\boldsymbol{\mu}_t - r_f \mathbf{1})^\top + c\,\mathbf{1}^\top$, the formula stated in the theorem.

Reading the row formula off element by element, $\Delta\mathbf{M}$ admits the additive split

$$\Delta\mathbf{M} \;=\; (\lambda - 1)\big(\mathbf{M} - r_f \mathbf{1}_T \mathbf{1}_N^\top\big) \;+\; c\,\mathbf{1}_T \mathbf{1}_N^\top, \tag{8}$$

which establishes the multiplicative-half / additive-half decomposition of the theorem statement. The first summand has rank $\operatorname{rank}(\mathbf{M} - r_f \mathbf{1}_T \mathbf{1}_N^\top)$, an integer between 1 and $\min(T, N)$. The second summand $c\,\mathbf{1}_T \mathbf{1}_N^\top$ is the outer product $c\,\mathbf{1}_T \otimes \mathbf{1}_N$ and so has rank one. Right-multiplying (8) by the invertible matrix $\mathbf{V}^{-1}/\gamma$ preserves rank in each summand, and rank subadditivity yields

$$\operatorname{rank}(\bar{\mathbf{W}}^{\text{post}} - \bar{\mathbf{W}}) \;\le\; \operatorname{rank}(\mathbf{M} - r_f \mathbf{1}_T \mathbf{1}_N^\top) + 1 \;\le\; \min(T, N) + 1,$$

the bound asserted by the theorem. Remark 3.1 characterises when this bound is achieved versus when the additive half is absorbed into the multiplicative half's column span. □

4 Corollaries

Three corollaries follow. §4.1 specialises Theorem 1 to the homogeneous-returns case and recovers the (C1)–(C3) spectral neutrality corollary of Frøseth (2026h). §4.2 derives the rank-scaling content of the theorem under ability-affine returns and identifies the absorption regime in which the additive half is subsumed by the multiplicative half's column span. §4.3 reads off the flow-only and stock-only sub-systems—recovering the differential-incidence result of Frøseth (2026c) and the drift-shift symmetry of Frøseth (2026g) respectively—and states the joint Pareto link to Frøseth (2026d).

4.1 Homogeneous-returns reduction

Corollary 1 (Homogeneous-returns reduction)**.** *Under the hypotheses of Theorem 1 with $\boldsymbol{\mu}_t \equiv \boldsymbol{\mu}$ for all t, the post-tax FOC shift is the rank-one matrix*

$$\bar{\mathbf{W}}^{\text{post}}_{\text{hom}} - \bar{\mathbf{W}}_{\text{hom}} \;=\; \mathbf{1}_T\, \boldsymbol{\delta}^\top, \qquad \boldsymbol{\delta} \;:=\; \tfrac{1}{\gamma}\, \mathbf{V}^{-1}\big((\lambda - 1)(\boldsymbol{\mu} - r_f \mathbf{1}) + c\, \mathbf{1}\big), \tag{9}$$

recovering the C1–C3 tax-neutrality corollary of Frøseth (2026h).

The proof appears in Section A.1.

4.2 Rank scaling under ability-affine returns

Corollary 2 (Rank scaling)**.** *Under the hypotheses of Theorem 1, suppose returns admit the ability-affine decomposition*

$$\boldsymbol{\mu}_t \;=\; \boldsymbol{\mu}_0 \;+\; \sum_{k=1}^{d} z_{t,k}\, \boldsymbol{\nu}_k, \tag{10}$$

with $\{\boldsymbol{\nu}_k\}_{k=1}^d \subset \mathbb{R}^N$ linearly independent and $\{\mathbf{1}_T, \mathbf{z}_1, \ldots, \mathbf{z}_d\} \subset \mathbb{R}^T$ jointly linearly independent (where $\mathbf{z}_k := (z_{1,k}, \ldots, z_{T,k})^\top$); in particular $d \le \min(T-1, N)$. Then the post-tax FOC shift has rank

$$\operatorname{rank}(\bar{\mathbf{W}}^{\text{post}} - \bar{\mathbf{W}}) \;=\; \begin{cases} \min(d+1,\, N) & \text{if } \boldsymbol{\mu}_0 \neq r_f \mathbf{1}, \\ \min(d,\, N) + 1 & \text{if } \boldsymbol{\mu}_0 = r_f \mathbf{1}. \end{cases} \tag{11}$$

The homogeneous-returns case $d = 0$ recovers Corollary 1 (rank one); a single-factor ability $d = 1$ gives rank two in either branch; the bound $\min(T, N) + 1$ of Theorem 1 is saturated only in the degenerate branch $\boldsymbol{\mu}_0 = r_f \mathbf{1}$ with $d = N$ and $T \ge N + 1$.

The proof appears in Section A.2.

Remark (Spectral mechanics of the rank-one bump)*.* When the rank-one bump materialises

($\boldsymbol{\mu}_0 = r_f\mathbf{1}$), the post-tax FOC shift is the sum of a rank-$\min(d, N)$ matrix (multiplicative-half post-$\mathbf{V}^{-1}$) and a rank-one matrix (additive-half $c\,\mathbf{1}_T(\mathbf{V}^{-1}\mathbf{1})^\top/\gamma$). The singular- value spectrum of the sum can be computed from the secular equation of Bunch et al. (1978) for the symmetric eigenproblem and from Bunch and Nielsen (1978) for its singular value decomposition (SVD) analogue: the post-update singular values interlace the pre-update ones, and the perturbation of the corresponding singular vectors is bounded by the canonical- angle estimates of Stewart (1973) via the multiple-eigenvalue condition number of Stewart and Zhang (1990). In the typical financial case $\boldsymbol{\mu}_0 \neq r_f\mathbf{1}$, the additive half is absorbed into the multiplicative half's column span and no rank-one update analysis is needed; the spectrum of the post-tax FOC shift coincides with the spectrum of $(\lambda - 1)(\mathbf{M} - r_f\mathbf{1}_T\mathbf{1}_N^\top)\mathbf{V}^{-1}/\gamma$.

Remark (Economic interpretation). The rank $\min(d + 1, N)$ counts the independent cross-sectional channels through which the (C1)–(C3) tax system distorts allocations under heterogeneous returns. Each ability factor contributes one flow-tax channel along $\mathbf{V}^{-1}\boldsymbol{\nu}_k$, contracting cross-agent dispersion along that direction by factor $1-\lambda$. A $(d+1)$-th channel along $\mathbf{V}^{-1}\mathbf{1}$ (the global-minimum-variance direction) can arise from either of two sources: the constant baseline $\boldsymbol{\mu}_0 - r_f\mathbf{1} \neq 0$ of the pre-tax expected returns, acting through the flow-tax multiplicative half, or the wealth tax acting through the additive half. In the financially typical case $\boldsymbol{\mu}_0 \neq r_f\mathbf{1}$ the two coincide and the wealth-tax contribution to the cross-sectional shift is statistically inseparable from the constant-baseline component of the flow-tax shift; only in the degenerate case $\boldsymbol{\mu}_0 = r_f\mathbf{1}$ does the wealth-tax channel appear as a genuinely independent direction. The homogeneous-returns baseline $d = 0$ collapses the entire cross-sectional shift onto $\mathbf{V}^{-1}\mathbf{1}$ and recovers the rank-one result of Corollary 1.

4.3 Connection to flow/stock incidence and joint Pareto

Corollary 3 (Flow/stock limits and joint Pareto link)**.** *Under the hypotheses of Theorem 1 with the ability-affine returns* (10)*:*

(a) Flow-only sub-system *($\tau_w = 0$). Then $\lambda < 1$ and $c = -(1-\tau_d)\tau_c r_f$. The multiplicative-half compresses the cross-row drift heterogeneity by $|\lambda - 1| > 0$: row t's shift in the tangency direction $\mathbf{V}^{-1}(\boldsymbol{\mu}_t - r_f\mathbf{1})$ scales by $\lambda - 1$, so the cross-agent dispersion of post-tax tangency portfolios is compressed relative to the pre-tax dispersion. This is the spectral-level statement of the differential-incidence result of Frøseth (2026c).*

(b) Stock-only sub-system *($\tau_c = \tau_d = 0$). Then $\lambda = 1$ and $c = -\tau_w a$. The multiplicative-half vanishes; only the rank-one additive-half $(c/\gamma)\,\mathbf{1}_T(\mathbf{V}^{-1}\mathbf{1})^\top$ remains. The shift is row-uniform: every agent's post-tax tangency portfolio is translated by the same $\boldsymbol{\delta}^{\text{stock}} = (c/\gamma)\,\mathbf{V}^{-1}\mathbf{1}$ along the global-minimum-variance direction, with cross-agent drift heterogeneity preserved (no compression). This recovers Frøseth (2026g)'s drift-shift symmetry, lifted to the heterogeneous-returns level.*

(c) Joint Pareto link. *Under the full (C1)–(C3) flow and stock tax system, the multiplicative-half compresses heterogeneity and the additive-half adds a row-uniform global-minimum-variance-direction shift. The post-tax fluctuation SDE inherits the same diffusion $\mathbf{V}$ as*

the pre-tax SDE (§2.3), so the long-run joint (log-wealth, ability) Pareto-exponent transformation derived in Frøseth (2026d, Theorem 1) follows from the deterministic FOC-shift formula of Theorem 1 combined with the fluctuation-SDE invariance argument of Frøseth (2026h).

The proof appears in Section A.3.

4.4 Differential wealth-tax assessment

Conditions (C1)–(C3) of Frøseth (2026b) include the uniform assessment requirement (C3), which the Norwegian dual income tax violates in practice through differential treatment of housing, listed equities, and bank deposits (Section 6.1). The following corollary lifts Theorem 1 from (C1)–(C3) to (C1) and (C2) alone, treating the assessment profile $\boldsymbol{\alpha} \in \mathbb{R}^N$ as a free vector. The multiplicative half is unchanged; only the rank-one additive half re-aligns to an asset-specific direction encoded in $\boldsymbol{\alpha}$.

Corollary 4 (Differential wealth-tax assessment)**.** *Under the hypotheses of Theorem 1 but with (C3) relaxed to allow asset-specific assessment $\alpha_i \neq \alpha_j$ across asset classes, the post-tax FOC shift on the allocation matrix is*

$$\bar{\mathbf{W}}^{\text{post}} - \bar{\mathbf{W}} \;=\; \tfrac{1}{\gamma}\big[(\lambda-1)(\mathbf{M} - r_f \mathbf{1}_T \mathbf{1}_N^\top) + \mathbf{1}_T\, \mathbf{c}^\top\big]\, \mathbf{V}^{-1}, \tag{12}$$

where $\mathbf{c} \in \mathbb{R}^N$ has components

$$c_j \;=\; -(1-\tau_d)\,\tau_c\, r_f \;-\; \tau_w\, \alpha_j. \tag{13}$$

The decomposition into a multiplicative half $(\lambda-1)(\mathbf{M}-r_f\mathbf{1}_T\mathbf{1}_N^\top)\mathbf{V}^{-1}/\gamma$ and a rank-one additive half $\mathbf{1}_T(\mathbf{V}^{-1}\mathbf{c})^\top/\gamma$ is inherited from Theorem 1, with the additive direction shifting from the global-minimum-variance direction $\mathbf{V}^{-1}\mathbf{1}$ to the assessment-weighted direction $\mathbf{V}^{-1}\mathbf{c}$. The rank bound $\min(T,N)+1$ and the ability-affine sharpening to $\min(d+1,N)$ of Corollary 2 carry over verbatim under the same absorption argument.

The proof appears in Section A.4.

Remark (Fluctuation-SDE invariance under broken (C3))*.* The Hessian of the per-row mean-variance loss with respect to $\mathbf{W}_{t\cdot}$ is $\gamma\mathbf{V}$ under both (C1)–(C3) and (C1)–(C2) with broken (C3): the asset-specific assessment enters only the linear term of the loss through $\mathbf{c}$, never the quadratic curvature. Consequently the fluctuation-SDE drift $-\eta\gamma\tilde{\mathbf{W}}_{t\cdot}\mathbf{V}$ established in Theorem 1 claim 3 is invariant under broken (C3) as well, and the cross- sectional bulk spectrum of the allocation matrix's fluctuations is preserved row-by-row. The agent-level wealth process, by contrast, acquires a portfolio-dependent wealth-tax cost $\tau_w\, \mathbf{W}_{t\cdot}\boldsymbol{\alpha}$ that has no analogue under (C3) and that does not factor out of the type-conditional Pareto- exponent computation the way the (C3)-uniform case does. The long-run distributional consequences of this coupling are discussed in Section 8.3.

4.5 Debt deductibility and the leveraged-FOC corollary

Corollary 4 treats $\boldsymbol{\alpha}$ as the bare assessment vector applied to gross asset values and assumes households hold unleveraged positions. Norwegian households leverage the discounted asset classes heavily through the residential mortgage market and through holding-company structures; the corresponding tax treatment of debt is intricate, and the institutional rules generate a clean algebraic carry-through of Corollary 4 only under a specific configuration we now isolate.

We refer to the post-2017 institutional configuration — proportional debt reduction (*gjeldsreduksjon*) on the wealth-tax base, and corporate-routed asset holding with uncapped interest deduction via corporate-loss carry-forward — as the *canonical case* for the leveraged analysis below. Three institutional features make this the clean configuration. First, Frøseth (2026a, §4.4) formalises the *gjeldsreduksjon* as a proportional rule under which the wealth-tax base attached to a leveraged position in asset class j is the discounted equity $\alpha_j(V_j - D_j)$ rather than the discounted gross value $\alpha_j V_j$. Second, when the asset is held through a holding company the personal interest deduction is replaced by corporate interest expense, deductible at τ_c against corporate income with unused deductions carried forward indefinitely as corporate losses; the deduction is therefore uncapped over time and operates at exactly the rate τ_c that the corporation pays on debt-financed gross income, restoring the (C1) symmetry on the debt margin. Third, the dividend stack with shielding (Section 2.3) under (C2) treats distributions of the after-corporate-tax retained amount uniformly. Together these three features collapse the leveraged tax stack onto the unleveraged structure of Corollary 4 up to a single transformation by the leverage matrix introduced below. Configurations outside this regime — directly held positions whose personal interest deduction is capped at gross capital income, and the pre-2017 full-deduction regime under which leveraged discounted assets shelter taxable wealth from other assets — are non-canonical in the sense that the symmetry argument breaks; we discuss these in Section 8.3.

Leveraged household balance sheet. Index the asset classes $j = 1, \ldots, N$ as in Section 2.1, with $\ell_j \in [0, 1)$ the loan-to-value (LTV) ratio on class j, V_j the gross asset value, $D_j = \ell_j V_j$ the associated debt, and $E_j = (1 - \ell_j)V_j$ the equity. The household's net wealth is total equity, $W = \sum_j E_j$. We work with *equity portfolio shares* $\theta_j := E_j/W$ throughout, so that $\sum_j \theta_j = 1$ and the gross asset value invested in class j is $V_j = \theta_j W/(1 - \ell_j)$. Collect the leverage ratios in the diagonal matrix

$$\Lambda \;:=\; \operatorname{diag}(1 - \ell_1, \ldots, 1 - \ell_N) \;\in\; \mathbb{R}^{N\times N}_{>0}, \tag{14}$$

so that $\Lambda \succ 0$ and $\Lambda = I$ recovers the unleveraged formulation of Section 2.1.

Post-tax excess return per unit of equity in class j. Under the canonical case the corporation borrows at the risk-free rate r_f (a normalisation; see Section 8.3 for the borrowing-spread case), pays corporate tax τ_c on $(\mu_j - r_f\ell_j)V_j$, distributes the after-corporate-tax retained amount $(1 - \tau_c)(\mu_j - r_f\ell_j)V_j$, and the household pays personal capital-income tax τ_k on the shielded portion $r_s(V_j - D_j)$ and dividend tax τ_d on the excess. Under (C1) ($\tau_k = \tau_c$) and (C2) ($r_s = r_f$), the household's post-tax excess return on asset class j over the risk-free asset, expressed per unit of equity E_j invested in class j, reduces to

$$R_j^{\text{ex,lev}}(\ell_j) \;=\; \frac{\lambda(\mu_j - r_f)}{1-\ell_j} \;+\; c_j, \qquad c_j \;=\; -(1-\tau_d)\tau_c r_f \;-\; \tau_w\,\alpha_j, \tag{15}$$

with λ and c as in (3). Two features deserve note. First, the gross-return premium $\lambda(\mu_j - r_f)$ is rescaled by the gearing factor $1/(1-\ell_j) \geq 1$: at LTV $\ell_j = 0.5$ the premium per unit of equity is doubled, at $\ell_j = 0.7$ it is tripled, and so on. Second, the additive component c_j is *unchanged* relative to the unleveraged Corollary 4. The wealth-tax wedge $\tau_w\,\alpha_j$ does not pick up a leverage factor: the *gjeldsreduksjon* reduces the wedge per unit of gross asset value to $\tau_w\,\alpha_j(1-\ell_j)$, but per unit of equity in class j the same factor cancels against $E_j = (1-\ell_j)V_j$ and the wedge returns to $\tau_w\,\alpha_j$. The leverage channel under the canonical case therefore acts through the gross-return premium alone.

Corollary 5 (Cross-sectional FOC shift under canonical leveraged routing)**.** *Under the hypotheses of Corollary 4 and the canonical case of Section 4.5 with row-uniform leverage matrix Λ, the post-tax FOC shift on the equity allocation matrix is*

$$\bar{\mathbf{W}}^{\text{post}} - \bar{\mathbf{W}} \;=\; \tfrac{1}{\gamma}\left[\mathbf{M}_{\text{centered}}\,\mathbf{V}^{-1}\,(\lambda\Lambda - I) \;+\; \mathbf{1}_T\,(\Lambda\,\mathbf{V}^{-1}\Lambda\,\mathbf{c})^{\top}\right], \tag{16}$$

where $\mathbf{M}_{\text{centered}} := \mathbf{M} - r_f \mathbf{1}_T \mathbf{1}_N^{\top}$ and $\mathbf{c}$ has components (15). The decomposition into a multiplicative half $\mathbf{M}_{\text{centered}}\mathbf{V}^{-1}(\lambda\Lambda - I)/\gamma$ and a rank-one additive half $\mathbf{1}_T(\Lambda\mathbf{V}^{-1}\Lambda\mathbf{c})^{\top}/\gamma$ is inherited from Theorem 1, with the unleveraged additive direction $\mathbf{V}^{-1}\mathbf{c}$ replaced by the leverage-weighted direction $\Lambda\mathbf{V}^{-1}\Lambda\mathbf{c}$. The rank bound $\min(T, N) + 1$ and the ability-affine sharpening to $\min(d + 1, N)$ of Corollary 2 carry through whenever $\lambda\Lambda - I$ has full rank, which holds for any $\Lambda \neq \lambda^{-1} I$.

The proof appears in Section A.5.

Remark (Fluctuation-SDE invariance and Theorem 2 under canonical leverage). The Hessian of the household's mean-variance loss with respect to the equity-share vector $\boldsymbol{\theta}_t$ is $\gamma\mathbf{V}^{\text{eq}} = \gamma\Lambda^{-1}\mathbf{V}\Lambda^{-1}$, positive-definite and row-wise constant under canonical leverage. The fluctuation-SDE drift $-\eta\gamma\tilde{\mathbf{W}}^{\text{post}}\mathbf{V}^{\text{eq}}$ is therefore invariant in form across the canonical-leveraged tax transformation, and Theorem 1 claim 3 carries through with $\mathbf{V}$ replaced by $\mathbf{V}^{\text{eq}}$ in the bulk-spectrum statement. The type-conditional Pareto exponent of Theorem 2 extends correspondingly: the optimal portfolio $\boldsymbol{\theta}^*(z)$ inherits the Λ-conjugated structure of Corollary 5, the wealth-process drift $m_X(z)$ involves the leveraged excess return (15), and the diffusion term picks up a Λ^{-1} factor on the risky portion. For homogeneous LTV $\ell_j = \ell$ the gearing acts as a uniform rescaling of the risky-asset risk premium and the closed form of Theorem 2 applies directly with $\mu_j - r_f$ replaced by $(\mu_j - r_f)/(1-\ell)$. The wealth-tax floor at zero on the household's aggregate taxable wealth, which lies outside this regime, is discussed in Section 8.3.

Wedge magnitude in the Norwegian application. Norwegian household balance sheets carry strongly asset-specific LTVs: primary housing typically has $\ell_H \approx 0.5$–0.7 through the residential mortgage market, listed equities and bank deposits are essentially unleveraged, and secondary housing sits between the two. The gearing factor $1/(1-\ell_j)$ on the gross-return premium is therefore strongly asset-specific in practice: housing's risk premium per unit of equity is amplified by a factor of two to three relative to the unleveraged baseline, while the equity and deposit premia are unchanged. The combined effect with the bare assessment profile of Table 2 is a per-unit-of-equity return profile that loads housing more heavily relative to the illustration in Section 6.1: the actual Norwegian household sits further toward the gearing-amplified housing direction in equity-allocation space than the bare-profile calibration suggests. Section 6.6 reports the quantitative consequence for the type-conditional Pareto exponent.

Scope: regimes outside the canonical case. Two institutional configurations sit outside the canonical case treated above. First, *directly held leveraged positions* (no intermediate corporate vehicle): the personal interest deduction is bounded by the household's gross capital income, and where this cap binds — as is common for households whose primary-housing mortgage interest exceeds dividend and deposit income — the effective deduction rate falls below τ_c and the (C1) symmetry on the debt margin breaks. The leveraged FOC shift then carries an additional asset-specific wedge whose magnitude depends on the realised cap; we treat this case as a corpus-wide revision item flagged in Section 8.3. Second, the *pre-2017 full deduction regime* ($\beta_j = 1$ uniformly under the notation of Frøseth (2026a, §4.4)): the wealth-tax contribution per unit of gross asset value, $\alpha_j - \ell_j$, can be negative when $\ell_j > \alpha_j$. The wealth-tax base is non-negative at the household level by construction, so a negative contribution from a leveraged discounted asset does not generate a negative tax payment; instead it *shelters* taxable wealth from other asset classes by offset, and for sufficiently leveraged configurations the household's aggregate base is reduced to the zero boundary. The closed-form Pareto exponent of Theorem 2 requires piecewise patching at the boundary in this regime; we likewise defer this case to Section 8.3.

Economic interpretation. Corollary 5 formalises a clean reading of the leverage channel under the canonical case: leverage acts on the FOC entirely through the gearing factor $1/(1-\ell_j)$ on the gross-return premium, leaving the wealth-tax wedge $\tau_w\,\alpha_j$ per unit of equity unchanged. The substantive consequence is that two households holding identical bare-$\boldsymbol{\alpha}$ exposures but different LTVs occupy different points in the leveraged FOC: the higher-LTV household's optimal equity share in discounted asset classes is amplified by the gearing factor, even though its wealth-tax wedge per unit of equity is identical. Cross-sectionally, the *within-asset* variation in LTV across households is a third source of dispersion in the post-tax allocation matrix, alongside return heterogeneity and assessment-fraction heterogeneity. The rank decomposition of Corollary 5 continues to constrain this dispersion: total rank is bounded by $\min(T, N) + 1$ and sharpens to $\min(d + 1, N)$ under the same ability-affine hypothesis, with the gearing factor entering the bound only through the diagonal $\lambda\Lambda - I$ rather than through the column-rank of $\mathbf{M}_{\text{centered}}$.

4.6 Progressive wealth taxation

The constant-τ_w assumption underlying Corollaries 4 and 5 is a simplification that the Norwegian implementation does not satisfy: since 2022 the wealth-tax schedule has been a two-bracket piecewise-constant function of net taxable wealth. The bracket-strand machinery of Frøseth (2026f, §4) gives closed-form steady-state distributions under such schedules in the homogeneous-returns case, and Frøseth (2026d, Cor. 2) extends the type-conditional Pareto formula of Theorem 2 to the bracket strand under (C3) and heterogeneous returns. We now extend Corollary 5 to piecewise-constant $\tau_w(x)$ schedules under broken (C3), recovering the bracket-conditional analogue of the leveraged FOC shift. The associated extension to the long-run Pareto exponent appears in Section 5.4.

Schedule and bracket assignment. Following Frøseth (2026f, §2), let the wealth-tax schedule be a piecewise-constant function of household net taxable wealth with K brackets $\{B_1, \dots, B_K\}$ partitioning $[0, \infty)$, and let $\tau_w^{(b)}$ denote the constant rate in bracket $b \in \{1, \dots, K\}$. The Norwegian post-2022 system corresponds to $K = 2$ with rates $(\tau_w^{(1)}, \tau_w^{(2)}) = (1.0\%, 1.1\%)$ and a threshold near NOK 20 M (Frøseth, 2026f, §5). Each agent t sits in exactly one bracket $b(t) \in \{1, \dots, K\}$ determined by their net taxable wealth at the FOC-equilibrium time; for the deterministic cross-sectional analysis we treat $b(t)$ as exogenously assigned. Define the bracket-indicator vectors $\mathbf{1}^{(b)} \in \{0,1\}^T$ with $\mathbf{1}_t^{(b)} = 1$ iff $b(t) = b$, satisfying $\sum_{b=1}^{K} \mathbf{1}^{(b)} = \mathbf{1}_T$.

Bracket-conditional shift constant. Under broken (C3) and canonical leverage with bracket-conditional $\tau_w^{(b)}$, the wealth-tax wedge per unit of equity in class j for an agent in bracket b is $\tau_w^{(b)} \alpha_j$, and the corresponding additive shift constant is

$$c_j^{(b)} \;=\; -(1-\tau_d)\,\tau_c\, r_f \;-\; \tau_w^{(b)}\,\alpha_j, \qquad \mathbf{c}^{(b)} \in \mathbb{R}^N. \tag{17}$$

The constant-τ_w case of Corollary 5 corresponds to $K = 1$ with $\mathbf{c}^{(1)} = \mathbf{c}$.

Corollary 6 (Bracket-conditional FOC shift under progressive wealth taxation)**.** *Under the hypotheses of Corollary 5 with the constant-τ_w assumption replaced by a K-bracket piecewise-constant schedule and bracket assignment $b(t)$ exogenous, the post-tax FOC shift on the equity allocation matrix is*

$$\bar{\mathbf{W}}^{\text{post}} - \bar{\mathbf{W}} \;=\; \tfrac{1}{\gamma}\,\mathbf{M}_{\text{centered}}\,\mathbf{V}^{-1}\,(\lambda\Lambda - I) \;+\; \tfrac{1}{\gamma}\sum_{b=1}^{K} \mathbf{1}^{(b)}\,\big(\Lambda\,\mathbf{V}^{-1}\Lambda\,\mathbf{c}^{(b)}\big)^{\top}. \tag{18}$$

The decomposition into a multiplicative half (rank $\le \min(T, N)$) and a sum of K rank-one bracket-conditional additive halves gives a total rank bound of $\min(T, N) + K$. Under the ability-affine hypothesis of Corollary 2 the multiplicative-half rank sharpens to $\min(d+1, N)$, giving total rank $\le \min(d+1, N) + K$. The $K = 1$ case recovers Corollary 5 verbatim; the $\Lambda = I$ case recovers the bracket-stratified version of Corollary 4.

The proof appears in Section A.6.

Remark (Bracket structure as a third source of dispersion)*.* The cross-sectional dispersion of the post-tax allocation matrix under progressive wealth taxation has three asset-side sources — heterogeneous returns, asset-specific assessment under broken (C3), and asset-specific leverage — plus one agent-side source: the bracket assignment $b(t)$. The first three enter the FOC shift through the multiplicative half and the Λ-conjugated additive direction; the bracket assignment enters through the rank-one indicator $\mathbf{1}^{(b)}$. The total rank bound $\min(T, N)+K$ is attained when the K bracket-indicator vectors and the column space of $\mathbf{M}_{\text{centered}}\mathbf{V}^{-1}(\lambda\Lambda - I)$ are mutually transversal in $\mathbb{R}^T$. For the Norwegian post-2022 case ($K = 2$) the bracket addition is one rank degree of freedom relative to Corollary 5, observable as a rank-three rather than rank-two leading singular structure of $\bar{\mathbf{W}}^{\text{post}} - \bar{\mathbf{W}}$ at $d = 1$.

Economic interpretation. Corollary 6 formalises a sharp empirical implication: under the actual Norwegian tax system, the cross-section of post-tax portfolio allocations carries an additional structural mode beyond those induced by heterogeneous returns and broken (C3). Households below the bracket threshold and households above it face different effective wealth-tax wedges ($\tau_w^{(1)}\boldsymbol{\alpha}$ versus $\tau_w^{(2)}\boldsymbol{\alpha}$), producing a step-function discontinuity in the per-agent FOC shift across the threshold. In matrix form this is the rank-one bracket-indicator outer product $\mathbf{1}^{(b)}(\Lambda\mathbf{V}^{-1}\Lambda\mathbf{c}^{(b)})^\top$, distinct in shape from both the multiplicative-half rescaling and the gjeldsreduksjon gearing channel. The bracket structure therefore enters the cross-sectional spectrum as a separate mode, additive to the multiplicative-and-leveraged structure of Corollary 5.

5 Long-run distributional consequences

Theorem 1 and Corollary 4 settle the deterministic side of the (C1)+(C2)+broken-(C3) story: the post-tax shift in the allocation matrix decomposes into a multiplicative half plus a rank-one additive half along the assessment-weighted direction $\mathbf{V}^{-1}\mathbf{c}$, and the fluctuation-SDE invariance of Theorem 1 claim 3 carries over (Section 4.4). What the deterministic story leaves open is the long-run distributional consequence: under (C3) the wealth-tax cost is identical across all agents and the tail of the post-tax wealth distribution responds only through the type-conditional drift; under broken (C3) the wealth-tax cost on agent t is $\tau_w \mathbf{W}_{t\cdot}\boldsymbol{\alpha}$, asset-portfolio dependent through the agent's holdings, with no analogue under (C3).

This section turns the deterministic FOC corollary into a closed-form type-conditional Pareto exponent. Section 5.1 bridges Corollary 4 to the joint Fokker–Planck framework of Frøseth (2026c,d); Section 5.2 states and proves the type-conditional Pareto-exponent transformation under broken (C3); Section 5.3 reads off three comparative-statics findings (non-monotonicity in the assessment wedge, the actual Norwegian profile sitting near the Pareto-exponent peak, and a sign-changing cross-type spread). The multi-asset Norwegian numerical illustration of these comparative statics is folded into the unified worked example of Section 6.6, alongside the FOC-level numerics of Theorem 1.

5.1 From Cor. 4 to the wealth process

For a type-z investor with optimal portfolio $\bar{\mathbf{W}}^*_{t\cdot}(z)$ given by Corollary 4 (specialised to a single representative risky asset for clarity; the multi-asset case follows by linearity in Section 5.2), let $w(z) := \bar{\mathbf{W}}^*_{t\cdot}(z)$ denote the post-tax fraction of wealth in the risky asset and $1 - w(z)$ the fraction in the risk-free anchor. Under (C1)+(C2) and the assessment profile (α_1, α_0), the type-conditional wealth process satisfies the geometric Brownian motion $\mathrm{d}X_t/X_t = m_X(z)\,\mathrm{d}t + \sqrt{\lambda}\,w(z)\,\sigma(z)\,\mathrm{d}W_t^{(z)}$ with per-dollar drift

$$m_X(z) \;=\; w(z)\,\big(\lambda(\mu(z) - r_f) + c_1\big) \;+\; (1 - w(z))\,(r_f - \tau_w\,\alpha_0), \tag{19}$$

where $c_1 = -(1 - \tau_d)\tau_c\, r_f - \tau_w\,\alpha_1$ as in Corollary 4, $\sigma(z)^2$ is the type-conditional asset-level return variance, and the $\sqrt{\lambda}$ diffusion factor inherits the flow-tax rescaling convention of Frøseth (2026f,d). Itô's formula on $Y_t = \log X_t$ yields the post-tax type-conditional log-wealth SDE

$$\mathrm{d}Y_t \;=\; v_{\mathrm{eff}}(z)\,\mathrm{d}t \;+\; \sqrt{2\,D_{\mathrm{eff}}(z)}\,\mathrm{d}W_t^{(z)}, \tag{20}$$

with effective drift and diffusion

$$v_{\mathrm{eff}}(z) = m_X(z) - \tfrac{1}{2}\,\lambda\,w(z)^2\sigma(z)^2, \tag{21}$$

$$D_{\mathrm{eff}}(z) = \tfrac{1}{2}\,\lambda\,w(z)^2\sigma(z)^2. \tag{22}$$

Under (C3) with $\alpha_1 = \alpha_0 = a$, Corollary 4 collapses to Corollary 1 and Frøseth (2026d) Lemma 1 (heterogeneous neutrality) restores the tax-blind optimal portfolio of Frøseth (2026b). Under broken (C3), (21)–(22) encode the portfolio–assessment coupling that drives the distributional results below.

5.2 Type-conditional Pareto exponent

Augmenting (20) with a type-conditional demographic source-sink at rate $\delta(z) > 0$, the joint stationary density $\pi(y, z)$ satisfies the joint Fokker–Planck equation of Frøseth (2026d, §4), and the type-conditional log-wealth marginal $\pi(y \mid z) \sim A(z)\,e^{-\alpha(z)\,y}$ for y large. The Pareto exponent is given by:

Theorem 2 (Type-conditional Pareto exponent under broken (C3))**.** *Under (C1)+(C2), broken (C3), heterogeneous returns parameterised by $(\mu(z), \sigma(z))$, and a type-conditional demographic source-sink at rate $\delta(z) > 0$, the upper-tail log-wealth marginal of the type-z investor's stationary wealth distribution decays at rate*

$$\alpha(z) \;=\; \frac{-\,v_{\mathrm{eff}}(z) \;+\; \sqrt{v_{\mathrm{eff}}(z)^2 + 4\,D_{\mathrm{eff}}(z)\,\delta(z)}}{2\,D_{\mathrm{eff}}(z)}, \tag{23}$$

with $v_{\mathrm{eff}}(z)$ and $D_{\mathrm{eff}}(z)$ given by (21)–(22) and $w(z)$ the post-tax optimal portfolio fraction of Corollary 4. The (C3) limit $\alpha_1 = \alpha_0$ recovers Frøseth (2026d, Theorem 1).

Proof. Substituting (20) into the joint Fokker–Planck equation of Frøseth (2026d, §4) and applying the eigenvalue argument for the upper-tail decay of a drift-diffusion process with a constant-rate source-sink (Frøseth, 2026d, §4.1) yields the closed form (23). Under (C3) with $\alpha_1 = \alpha_0$, Corollary 4 reduces to Corollary 1 and $w(z)$ aligns with the tax-blind optimal portfolio of Frøseth (2026b); substituting into (21) and (22) recovers $v_{\mathrm{eff}}(z) = \lambda\,m_0(z) - r$ and $D_{\mathrm{eff}}(z) = \lambda\,\sigma(z)^2/2$

of Frøseth (2026d, §3) with $r = \tau_w\, a$, verbatim. $\square$

Remark (Economic interpretation). The type-conditional Pareto exponent $\alpha(z)$ measures the upper-tail decay rate of the type-z stationary log-wealth density: a larger $\alpha(z)$ means a lighter tail and less wealth concentration for that type. Under (C3) the wealth-tax channel collapses to a uniform constant, and $\alpha(z)$ responds only to the type-conditional risk premium $\mu(z) - r_f$, the flow-tax retention λ, and demographic turnover $\delta(z)$. Under broken (C3) the agent's optimal portfolio $w(z)$ responds to the asset-class assessment profile $\boldsymbol{\alpha}$ via Corollary 4, and the wealth-process volatility $\lambda w(z)^2\sigma(z)^2$ becomes a function of both the type's risk premium and the assessment wedge $\alpha_1 - \alpha_0$. This couples the external $\boldsymbol{\alpha}$-profile to the long-run wealth distribution in a way that breaks symmetry across types: the same wedge produces lighter tails for some types and heavier tails for others (Section 5.3). The substantive distributional content under broken (C3) lives in this coupling and has no analogue at the (C3)-clean level.

Remark (Multi-asset extension). The single-risky-asset specialisation above is not essential. For N risky assets with covariance $\mathbf{V}$, type-conditional drift $\boldsymbol{\mu}(z)$, and assessment vector $\boldsymbol{\alpha}$, Corollary 4 gives the optimal portfolio $\bar{\mathbf{W}}^*_{t\cdot}(z) = \frac{1}{\gamma}\mathbf{V}^{-1}(\lambda(\boldsymbol{\mu}(z) - r_f\mathbf{1}) + \mathbf{c})$ with $\mathbf{c}$ asset-specific. The post-tax type-conditional log-wealth SDE then has effective drift and diffusion $v_{\text{eff}}(z) = \bar{\mathbf{W}}^{*\top}(z)(\lambda(\boldsymbol{\mu}(z) - r_f\mathbf{1}) + \mathbf{c}) + (1 - \bar{\mathbf{W}}^{*\top}(z)\mathbf{1})(r_f - \tau_w\alpha_0) - \frac{1}{2}\,\lambda\,\bar{\mathbf{W}}^{*\top}(z)\mathbf{V}\bar{\mathbf{W}}^*(z)$ and $D_{\text{eff}}(z) = \frac{1}{2}\,\lambda\,\bar{\mathbf{W}}^{*\top}(z)\mathbf{V}\bar{\mathbf{W}}^*(z)$. The Pareto formula (23) carries through verbatim with these multi-asset $v_{\text{eff}}(z)$ and $D_{\text{eff}}(z)$. The (C3) limit $\boldsymbol{\alpha} = a\,\mathbf{1}$ again recovers Frøseth (2026d, Theorem 1).

Remark (Mean-field counterpart). The heterogeneous-returns wealth process behind Theorem 2 has a mean-field counterpart in Bernard et al. (2026), who solve the model of quenched heterogeneous growth rates coupled by a mean-field redistribution term and identify a phase transition between localised (condensed) and delocalised wealth distributions at a critical redistribution intensity, with the tail exponent in the partially localised phase governed by the growth-rate *variance*. In the present framework the cross-type dispersion of $v_{\text{eff}}(z)$ and $D_{\text{eff}}(z)$ plays that role: the sign-changing cross-type spread of Section 5.3 is a tax-design-induced reshaping of the same dispersion that drives their transition, reached here through the joint Fokker–Planck route of Frøseth (2026d) rather than through the mean-field limit.

5.3 Comparative statics

Three comparative-statics readings of Theorem 2 give the substantive new content beyond Frøseth (2026d, Theorem 1)'s (C3)-clean baseline.

Non-monotonicity in the assessment wedge. Hold $(\boldsymbol{\mu}(z), \mathbf{V}, r_f, \gamma, \tau_c, \tau_d, \tau_w, \alpha_0)$ fixed and vary the risky-asset assessment α_1 (single-asset case) or α_j for one asset (multi-asset). The Pareto exponent $\alpha(z)$ is a non-monotonic function of the assessment wedge $\alpha_1 - \alpha_0$. At small wedges the agent's optimal portfolio fraction $w(z)$ has the same sign as the after-tax expected risk premium and is bounded away from zero; as α_1 rises, the wealth-tax wedge $-\tau_w(\alpha_1 - \alpha_0)$ adds to the risk premium's drag, $w(z)$ shrinks toward zero, and the diffusion $D_{\text{eff}}(z) = \lambda w(z)^2\sigma(z)^2/2$

falls to its minimum. The Pareto exponent $\alpha(z)$ peaks at the α_1 where $w(z) = 0$ — the agent's portfolio is exactly indifferent and the wealth-process volatility vanishes — and falls as $|w(z)|$ grows on either side. Economically, the agent uses portfolio weight to absorb the wealth-tax wedge; the wedge value that drives $w(z)$ to zero deterministicises the wealth process and gives the lightest possible upper tail. Wedges on either side reactivate volatility and thicken the tail.

The Norwegian profile sits near the peak. At calibrated Norwegian parameters ($\tau_c = 0.22$, $\tau_d = 0.378$, $\tau_w = 0.011$, $r_f = 0.03$, $\gamma = 10$), the assessment wedge that sets $w(z) = 0$ is approximately $\alpha_1 - \alpha_0 \approx 1$ for the equity-like risky asset and $\alpha_h \approx 0.25$ for housing. Norway's actual housing assessment of $\alpha_h = 0.25$ from Table 2 sits exactly at the housing peak; the equity assessment of $\alpha_e = 0.80$ is close to the equity peak. In the wedge-magnitude sweep of Figure 2(a) — scaling the entire profile from $\boldsymbol{\alpha} = \mathbf{1}$ (full (C3)) to $\boldsymbol{\alpha} = \boldsymbol{\alpha}_{\text{Norway}}$ via $\boldsymbol{\alpha}(\xi) = \mathbf{1} + \xi(\boldsymbol{\alpha}_{\text{Norway}} - \mathbf{1})$ for $\xi \in [0, 1.5]$ — the Pareto exponent peaks at $\xi \approx 1$, the actual Norwegian profile. Economically, the housing assessment of 0.25 that Norway settled on in 2010 as a political compromise on owner-occupied housing (Frøseth, 2026a, §5.1) happens to coincide with the wedge value that minimises long-run wealth concentration at the top of the distribution at representative ability. Whether by intent or by accident, Norway's anisotropic-by-housing assessment is approximately the optimum from a Pareto-tail-suppression perspective.

Sign-changing cross-type spread. Define the cross-regime Pareto spread

$$\Delta\alpha(z) \;:=\; \alpha(z; \text{broken (C3)}) - \alpha(z; \text{(C3)}). \tag{24}$$

Sweeping the type-conditional return level $\boldsymbol{\mu}(z)$ across μ-scale factors representing low- to high-ability investors, $\Delta\alpha(z)$ is positive at low $\boldsymbol{\mu}(z)$ and negative at high $\boldsymbol{\mu}(z)$. Low-ability agents see a lighter post-tax tail under broken (C3): their portfolio moves closer to the wealth-tax-optimal point, reducing volatility; high-ability agents see a heavier post-tax tail because their portfolio moves further away. The sign-change occurs at an intermediate $\boldsymbol{\mu}(z)$ that depends on the calibration but exists generically (Figure 2(b)). Economically, broken-(C3) assessment is regressive in the tail mechanism: it makes the long-run wealth distribution *lighter* at the bottom (where low-ability investors hold cash-like high-α assets and bear the brunt of uniform wealth tax under the (C3) counterfactual) and *heavier* at the top (where high-ability investors load up on equity-like medium-α assets and the assessment wedge gives them more risky exposure than under (C3)). Anisotropic Norwegian assessment compresses the bottom of the wealth distribution and concentrates the top relative to (C3)-uniform — a finding that runs counter to a naive "lower-assessment-equals-thinner-tail" reading.

5.4 Type-conditional Pareto exponent under progressive wealth taxation

Theorem 2 characterises the type-conditional upper-tail decay rate under constant τ_w. Norway's actual two-bracket schedule (Section 4.6) departs from this by setting a higher rate $\tau_w^{(2)}$ above a wealth threshold. The next proposition extends Theorem 2 to piecewise-constant $\tau_w(x)$ schedules, exploiting the fact that the upper-tail decay is governed by the topmost bracket alone.

Proposition 1 (Type-conditional Pareto exponent under progressive wealth taxation)**.** *Under the hypotheses of Theorem 2 with the constant-τ_w assumption replaced by a K-bracket piecewise-constant schedule $\tau_w(x)$ (Section 4.6) having upper-bracket rate $\tau_w^{(K)}$, the type-conditional Pareto exponent of the upper-tail log-wealth distribution is given by (23) with τ_w replaced throughout by $\tau_w^{(K)}$. The upper-tail exponent is therefore independent of both the threshold structure $\{X_b\}_{b=1}^{K-1}$ and the lower bracket rates $\{\tau_w^{(b)}\}_{b=1}^{K-1}$; only the topmost bracket rate enters.*

Proof. The argument of Theorem 2 characterises the upper-tail decay of the stationary type-conditional log-wealth distribution by solving the type-conditional Fokker–Planck eigenvalue ODE on the upper bracket of the schedule, where $\tau_w(x) = \tau_w^{(K)}$ is constant. The eigenvalue ODE on this bracket is identical in form to the constant-τ_w case with $\tau_w \to \tau_w^{(K)}$, so the closed form (23) applies verbatim with this substitution. The bracket-strand recovery argument of Frøseth (2026d, Cor. 2) establishes that lower brackets shape the steady-state distribution within those brackets but do not enter the upper-tail decay rate. $\square$

Remark (Implication for Norwegian comparative statics)*.* For the Norwegian post-2022 system with $(\tau_w^{(1)}, \tau_w^{(2)}) = (1.0\%, 1.1\%)$, the upper-tail $\alpha(z)$ depends only on $\tau_w^{(2)} = 1.1\%$ — exactly the constant rate used in the §6 calibration. The numerical illustration of Section 6.6 therefore reads unchanged under Proposition 1, with the constant rate reinterpreted as the upper-bracket rate. The substantive new content of Proposition 1 is the *policy invariance* statement: changing $\tau_w^{(1)}$ or the bracket threshold leaves the upper-tail Pareto exponent unchanged. Lower-bracket policy levers shape the middle of the wealth distribution, not its upper tail.

6 Norwegian worked example

This section illustrates the four channels of the framework on a single Norwegian calibration: heterogeneous returns (Theorem 1, Corollary 2), differential wealth-tax assessment (Corollary 4), canonical leverage under the post-2017 *gjeldsreduksjon* and corporate routing (Corollary 5), and progressive bracket schedules (Corollary 6, Proposition 1). The institutional anchor is the post-2017 Norwegian dual income tax: conditions (C1) and (C2) hold by design under the *fritaksmetoden* participation exemption; the empirical deviations are non-uniform assessment α_i across asset classes (Table 2), heavily-leveraged primary housing through the residential mortgage market, and the two-bracket schedule introduced in 2022. The FOC-level illustration in Sections 6.1 to 6.4 adopts the (C3)-counterfactual uniform assessment $a = 1$ to exercise Theorem 1 cleanly; Section 6.5 and Figure 1(c) read the actual asymmetric Norwegian profile through Corollary 4; Section 6.6 extends to the long-run Pareto exponent of Theorem 2 and reports the leveraged Norwegian Pareto value under Corollary 5; and the closing paragraph of Section 6.6 reads the same calibration under the post-2022 two-bracket schedule via Proposition 1. The ability-stratified return structure of Frøseth (2026d) supplies the row heterogeneity needed throughout.

6.1 Calibration source

Norwegian institutional design under the post-2017 fritaksmetoden satisfies (C1) by setting the capital-income tax rate equal to the corporate tax rate, $\tau_k = \tau_c = 22\%$, and (C2) by setting the shielding rate equal to the risk-free rate, $r_s = r_f$. The remaining (C3) condition—uniform assessment $\alpha_i = \alpha$ across asset classes—is the binding deviation in practice. The 2026 Norwegian assessment fractions are reproduced in Table 2 from Frøseth (2026a, Table 2): bank deposits and secondary housing are assessed at full market value ($\alpha = 1$), listed shares and commercial property at a 20% discount ($\alpha = 0.80$), and primary housing under NOK 10m at a 75% discount ($\alpha = 0.25$). The discounts are large and asset-class-specific: Frøseth (2026a) estimates the resulting portfolio tilt, computed via the proportional-tax FOC, at over 18 percentage points toward primary housing for a representative investor at $\tau_w = 1.0\%$ and $\gamma = 4$. Frøseth (2026b) characterises the same asymmetry as the dominant flow-tax-equivalent distortion under the Norwegian system.

Table 2: Norwegian wealth-tax assessment fractions (2026), reproduced from Frøseth (2026a, Table 2). Source: Skatteetaten (2026).

Asset class	Assessment fraction α_i	Effective discount
Bank deposits	1.00	0%
Secondary housing	1.00	0%
Listed shares	0.80	20%
Unlisted shares	0.80	20%
Commercial property	0.80	20%
Holiday homes	0.30	70%
Primary housing ($\leq$ NOK 10m)	0.25	75%
Primary housing ($>$ NOK 10m)	0.70	30%

Ability stratification follows the type-conditional drift framework of Frøseth (2026d): each household's expected-return vector $\boldsymbol{\mu}(z_t)$ inherits a small-d factor structure across the cross-section, with z_t the household's ability label.

6.2 Combined-tax parameters and derived scalars

The parameter values used here, drawn from the post-2017 Norwegian dual income tax and the (C3)-counterfactual uniform assessment, are

$$\tau_c = 0.22, \quad \tau_d = 0.378, \quad \tau_w = 0.011, \quad a = 1, \quad r_f = 0.03, \tag{25}$$

with risk-aversion coefficient $\gamma = 3$ taken as a representative mean-variance calibration. Substituting into (3) gives the combined-tax derived scalars

$$\lambda = (1 - \tau_c)(1 - \tau_d) = 0.485, \qquad c = -(1 - \tau_d)\tau_c r_f - \tau_w\, a = -0.0151, \tag{26}$$

and the per-row compression and translation factors

$$|\lambda - 1| = 0.515, \qquad c/\gamma = -0.00504. \tag{27}$$

The flow-tax channel multiplies each household's pre-tax excess-return signal by λ, just under halving it; the wealth-tax channel adds a row-uniform translation along the $\mathbf{V}^{-1}\mathbf{1}$ direction (the global-minimum-variance portfolio direction up to a normalising constant) with coefficient $c/\gamma \approx -0.005$, an order of magnitude smaller per-row than the multiplicative-half perturbation on a typical excess-return signal in the 2–6 percentage-point range.

Economic interpretation of (λ, c). The factor $\lambda \approx 0.485$ is the post-tax retention rate on the risk premium of every risky asset under the Norwegian dual income tax: just under half of a household's pre-tax expected excess return is retained net of corporate and dividend taxes. In the Domar–Musgrave reading, the Norwegian state is a silent partner collecting roughly 51.5% of the pre-tax risk premium and bearing the same fraction of the loss. The wealth-tax channel $c/\gamma \approx -0.005$ acts on *position* rather than *return* and is quantitatively dominated by the multiplicative compression at this calibration: the wealth-tax rate of 1.1% contributes an order of magnitude less per-row distortion than the silent-partner cut. The two channels are qualitatively distinct in their cross-sectional incidence — the multiplicative-half compresses ability-driven dispersion in the risk-premium signal, while the additive-half translates every household uniformly toward the GMV portfolio — and their separate identification (Corollary 3) is the empirical content of distinguishing flow-tax effects from stock-tax effects in register-data spectra.

6.3 Ability stratification and the cross-sectional panel

For the spectral computation we work with a $T = 8$, $N = 6$ panel: eight households indexed by ability label z_t, allocating across six risky assets. Following the ability-affine structure of Corollary 2 with $d = 2$ ability factors,

$$\boldsymbol{\mu}_t \;=\; \boldsymbol{\mu}_0 \;+\; z_{t,1}\boldsymbol{\nu}_1 \;+\; z_{t,2}\boldsymbol{\nu}_2, \tag{28}$$

the baseline drift $\boldsymbol{\mu}_0 \in \mathbb{R}^N$ is centred at 6% with cross-asset standard deviation 5 pp; the two factor-loading vectors $\boldsymbol{\nu}_1, \boldsymbol{\nu}_2 \in \mathbb{R}^N$ have standard deviation 2 pp; and the household ability scores $z_{t,k}$ are drawn from a standard normal. The asset covariance $\mathbf{V} \in \mathbb{R}^{N\times N}$ is a positive-definite random draw representing the cross-asset return-fluctuation structure. The baseline drift $\boldsymbol{\mu}_0 \neq r_f\mathbf{1}$ throughout, placing the panel in the typical financial-parametrisation regime of Remark 3.1 where the additive-half is absorbed into the multiplicative-half's column span.

The choice $d = 2$ is parsimonious. Fagereng et al. (2020) document that two broad channels—scale (top-decile vs the rest) and skill (within-decile return persistence)—explain most of the cross-decile variance in returns in the Norwegian register data, and Frøseth (2026d) use a similarly low-dimensional ability specification in its closed-form Pareto-exponent analysis. The qualitative content of the rank-scaling result extends to any d through Corollary 2.

6.4 Allocation-matrix shift, singular spectrum, and rank check

Computing $\Delta\mathbf{M}$ via Theorem 1 and the post-tax allocation-matrix shift $\bar{\mathbf{W}}^{\text{post}} - \bar{\mathbf{W}} = \frac{1}{\gamma}\Delta\mathbf{M}\mathbf{V}^{-1}$ on the panel of §6.3 yields three nonzero singular values:

$$\sigma_1 = 0.0525, \qquad \sigma_2 = 0.00924, \qquad \sigma_3 = 0.00497. \tag{29}$$

The singular indices $i = 4, 5, 6$ are zero to machine precision. Figure 1(a) plots these alongside the analogous spectra of the multiplicative-half and additive-half contributions taken in isolation: the multiplicative-half also has exactly three nonzero singular values, the additive-half has exactly one (consistent with its rank-1 outer-product form $(c/\gamma)\,\mathbf{1}_T(\mathbf{V}^{-1}\mathbf{1})^\top$), and the two halves contribute comparable spectral mass on this calibration despite the order-of-magnitude gap in their per-row contributions established in §6.2.

The rank check confirms the absorption regime of Remark 3.1: the multiplicative-half has rank 3, the additive-half has rank 1, and their sum has total rank 3, not 4. The additive-half's column space is $\operatorname{span}(\mathbf{1}_T)$, which already lies in the multiplicative-half's column space $\operatorname{span}(\mathbf{1}_T, \mathbf{z}_1, \mathbf{z}_2)$ because $\boldsymbol{\mu}_0 \neq r_f\mathbf{1}$. The additive-half's contribution therefore overlays the existing column space without expanding it, modifying the singular values of the sum without producing a fourth principal direction. The homogeneous-returns special case ($d = 0$) recovers a rank-1 shift exactly, as Corollary 1 predicts and as the spectral neutrality result of Frøseth (2026h) requires.

Figure 1(b) traces the rank scaling across $d \in \{0, 1, \ldots, 6\}$. The post-tax FOC shift's rank rises linearly from 1 at $d = 0$ to $\min(T, N) = 6$ at $d = 5$, where it saturates. The Norwegian case $d = 2$ corresponds to a rank-3 cross-sectional perturbation: a tripling of the spectral footprint of the (C1)–(C3) tax stack relative to the homogeneous-returns baseline.

Economic reading of the rank-three structure. The three nonzero singular values correspond to three independent cross-sectional distortion modes that the Norwegian (C1)–(C3) tax stack induces on household portfolios. The leading mode (σ_1) is the row-uniform shift toward the average pre-tax drift $\boldsymbol{\mu}_0$, common across households of all ability levels and inherited from the homogeneous-returns spectral neutrality of Frøseth (2026h). The second and third modes (σ_2, σ_3) are ability-conditional: they capture the tax stack's differential treatment of high- and low-ability households along the two ability factors. Empirically, this predicts that the post-tax cross-section of Norwegian household portfolios is governed by three principal modes — a level shift common to all households, plus two ability-stratified rotations that distinguish the top of the return distribution from the bottom — rather than the single uniform direction that a representative-investor model would predict. A reduced-form regression treating the cross-section as homogeneous in returns would conflate the three modes and miss two of them entirely.

6.5 Reading the spectrum: observable consequences

The substantive reading is that the Norwegian (C1)–(C3) tax system, even in its institutionally clean limit (uniform assessment, fritaksmetoden compliance), induces three independent cross-

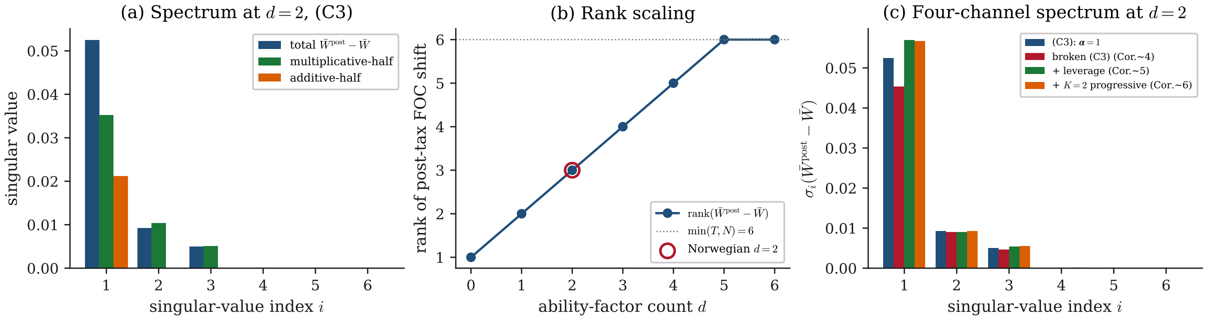

Figure 1: Spectrum and rank of the post-tax FOC shift on the illustrative Norwegian panel of §6. *Panel (a):* singular values of the total post-tax FOC shift $\bar{\mathbf{W}}^{\text{post}} - \bar{\mathbf{W}}$ alongside its multiplicative-half and additive-half decomposition (Theorem 1) for the $d = 2$ ability stratification with $T = 8$, $N = 6$ under (C1)–(C3) with the $a = 1$ counterfactual. Three nonzero singular values; the additive-half is rank-1 and is absorbed into the multiplicative-half's column span (no fourth principal direction; Remark 3.1). *Panel (b):* rank of the post-tax FOC shift as a function of the ability-factor count d. The rank rises linearly from $d = 0$ (homogeneous baseline, rank 1, recovering Frøseth (2026h)) up to $d = 5$, saturating at $\min(T, N) = 6$. The Norwegian case $d = 2$ (red marker) is a rank-3 cross-sectional perturbation. *Panel (c):* four-channel comparison of the total post-tax FOC shift's singular spectrum at $d = 2$. (i) (C3) baseline ($\boldsymbol{\alpha} = \mathbf{1}$, the §5 counterfactual). (ii) Broken (C3) (Corollary 4) at the actual Norwegian assessment profile $\boldsymbol{\alpha} = (0.25, 0.25, 0.80, 0.80, 1, 1)^\top$. (iii) Broken (C3) plus canonical leverage (Corollary 5) at representative housing LTV $\ell_H = 0.6$ on the two housing slots and zero on equities and deposits. (iv) Broken (C3) plus leverage plus the post-2022 $K = 2$ progressive schedule (Corollary 6) at rates $(\tau_w^{(1)}, \tau_w^{(2)}) = (1.0\%, 1.1\%)$ with half the agents in each bracket. Cases (i)–(iii) have rank 3; case (iv) has rank 4, the bracket-stratification rank addition predicted by Corollary 6.

sectional distortion directions in the post-tax allocation matrix—not the single direction that the homogeneous-returns reading of Frøseth (2026h) would predict. Two of the three directions are pure ability-factor distortions: their existence is conditional on the empirical fact of return heterogeneity, and they collapse onto the homogeneous-baseline direction if return dispersion across households is shut down (the limit $\boldsymbol{\nu}_1, \boldsymbol{\nu}_2 \to 0$, equivalently $d \to 0$). The first and largest direction is the row-uniform $\boldsymbol{\mu}_0$-aligned shift that survives in the homogeneous limit; it carries the bulk of the spectral mass at this calibration ($\sigma_1/\sigma_2 \approx 5.7$).

The connection to the long-run wealth distribution is mediated by the joint Fokker–Planck dynamics: each rank in the post-tax FOC shift contributes a corresponding direction in the deterministic-equilibrium drift of the (log-wealth, ability) joint density, and the closed-form Pareto-exponent transformation of Frøseth (2026d, Theorem 1) reads off the resulting tail-exponent change from this drift shift combined with the diffusion-invariance of the post-tax SDE (§2.3). The differential-incidence reading of Frøseth (2026c)—flow taxes compress and stock taxes preserve heterogeneity at the marginal-distribution level—therefore extends to the cross-sectional spectral level with explicit rank scaling: the combined flow and stock tax system distorts portfolio composition along multiple spectral channels simultaneously, with each ability factor opening a new channel.

The (C3)-counterfactual override $a = 1$ adopted in Section 6.1 for clean exposition can be lifted via Corollary 4. Substituting the actual Norwegian profile $\boldsymbol{\alpha} = (0.25, 0.25, 0.80, 0.80, 1, 1)^\top$ (two panel slots each for housing, listed equities, and bank deposits) into (12) leaves the multiplicative half unchanged but rotates the rank-one additive-half direction from $\mathbf{V}^{-1}\mathbf{1}$ to $\mathbf{V}^{-1}\mathbf{c}$, an assessment-weighted combination tilted toward the heavily-assessed asset classes. The rank scaling is unchanged at $\min(d+1, N) = 3$ at $d = 2$, and the total post-tax FOC shift's three nonzero singular values shift from $(\sigma_1, \sigma_2, \sigma_3) = (0.0525, 0.00924, 0.00497)$ under (C3) to $(0.0454, 0.00902, 0.00465)$ under broken (C3) — a 14% compression of the leading singular value and modest shifts of the trailing two (Figure 1(c)). The empirical reading carries through: the predicted post-tax FOC shift in register data still has rank three at $d = 2$, but the second principal direction is no longer the global-minimum-variance portfolio; it is the assessment-weighted direction $\mathbf{V}^{-1}\mathbf{c}$, identifiable as a direct diagnostic of which asset classes bear the wealth-tax burden disproportionately.

The Norwegian housing tilt as a spectral signature. The shift of the additive direction from $\mathbf{V}^{-1}\mathbf{1}$ to $\mathbf{V}^{-1}\mathbf{c}$ is the cross-sectional fingerprint of the Norwegian housing tilt. Primary housing's $\alpha = 0.25$ assessment discount creates a 75% valuation gap relative to bank deposits at $\alpha = 1.00$, and the post-tax FOC's additive direction loads disproportionately on housing as a result. In Norwegian register data, this tilt manifests as the well-documented overweighting of primary housing in upper-decile portfolios (Fagereng et al., 2024); the spectral signature predicted here is the sharper statement that the *cross-sectional* direction of the post-tax allocation shift is asset-specific and identifiable from register-data SVDs. Households whose pre-tax allocation already overweights housing absorb the post-tax distortion through their existing exposure; households heavy in fully-assessed assets (deposits, secondary housing) bear more of the post-tax adjustment burden through portfolio rebalancing. The 14% compression of the leading singular value under broken (C3) relative to (C3) is the empirical analogue of this asymmetric incidence at the spectral level.

6.6 Long-run Pareto-exponent illustration

A multi-asset numerical illustration confirms the comparative statics. The setup uses three risky asset classes (housing, listed equities, commercial property) at Table 2 assessment fractions $(\alpha_h, \alpha_e, \alpha_p) = (0.25, 0.80, 0.80)$, representative pre-tax returns $\boldsymbol{\mu} = (4.5\%, 7.5\%, 6.0\%)$, volatilities $(\sigma_h, \sigma_e, \sigma_p) = (10\%, 20\%, 15\%)$, the correlation profile of Section 6.3 (housing–equities 0.20, housing–property 0.30, equities–property 0.40), and type-conditional demographic turnover $\delta(z) = 6.2\%$ following the calibration of Frøseth (2026d, §4.1).[1]

Figure 2 reports the two headline findings from this calibration. Panel (a) traces the Pareto exponent as the wedge-magnitude parameter ξ varies from 0 (full (C3) baseline) to 1.5 (overshooting the actual Norwegian wedge). The unleveraged exponent peaks near $\xi = 1$ at $\alpha(z) \approx 2.052$, with the (C3) baseline at $\alpha(z) \approx 2.007$ and the overshooting case at $\alpha(z) \approx 2.039$ — the actual Norwegian bare profile is the empirical analogue of the wedge-magnitude that maximises the

[1] The risk-aversion parameter $\gamma = 10$ here differs from the $\gamma = 3$ of Section 6.2: the Pareto-exponent computation requires a realistic optimal-portfolio composition (modest risky weights, not extreme leverage), which the higher risk-aversion calibration delivers. The deterministic FOC shift of Theorem 1 is robust to γ; the long-run distributional analysis here is not.

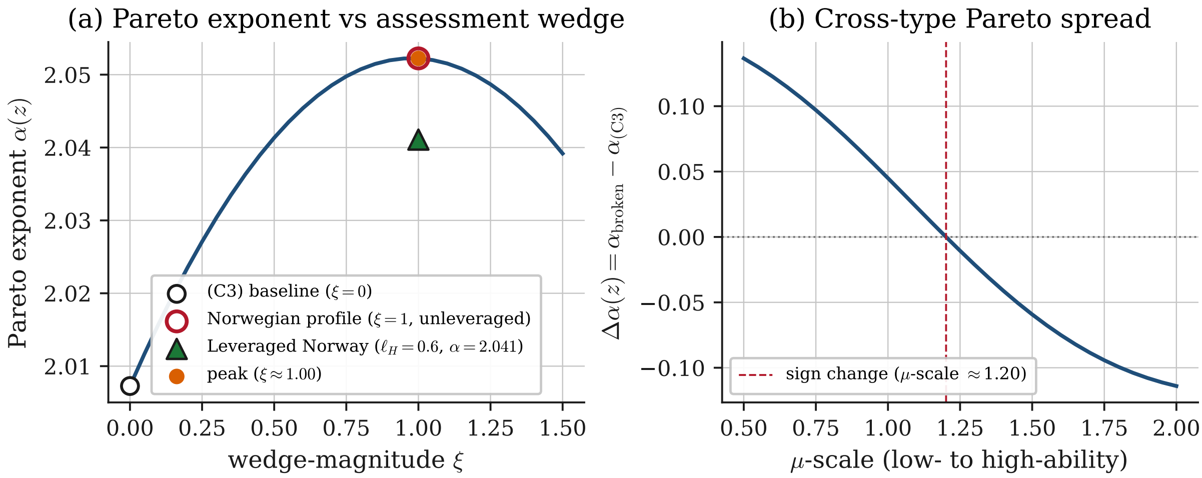

Figure 2: Long-run distributional consequences under broken (C3). *Panel (a):* Pareto exponent $\alpha(z)$ as a function of the wedge-magnitude parameter ξ, where $\boldsymbol{\alpha}(\xi) = \mathbf{1} + \xi(\boldsymbol{\alpha}_{\text{Norway}} - \mathbf{1})$. The unleveraged exponent peaks near $\xi = 1$ — the actual Norwegian profile — with both the (C3) baseline ($\xi = 0$) and the overshooting regime ($\xi > 1$) producing heavier upper tails. The triangle marker at $\xi = 1$ reports the leveraged Norwegian Pareto exponent under canonical leverage (Corollary 5) at representative housing LTV $\ell_H = 0.6$. *Panel (b):* cross-type spread $\alpha(z; \text{broken}) - \alpha(z; (\text{C3}))$ as a function of the μ-scale factor parameterising low- to high-ability investors. The spread is positive at low ability (broken (C3) compresses the bottom), crosses zero at intermediate ability, and is negative at high ability (broken (C3) concentrates the top).

upper-tail Pareto exponent at a representative-ability level. Adding canonical leverage (Corollary 5) at the representative housing LTV $\ell_H = 0.6$ shifts the Norwegian-profile exponent to $\alpha(z) \approx 2.041$ — a downward shift of about 0.011 relative to the unleveraged case, indicating that the gearing amplification of the gross-return premium produces a marginally heavier upper tail through the equity-share rescaling alone, with the wealth-tax wedge per unit of equity unchanged (Section 4.5). Economically, the leverage shift is identifiable from the wealth-tax shift on a fundamental specification ground: leverage acts entirely on the gross-return premium ($1/(1 - \ell_j)$ gearing factor), while the wealth-tax wedge acts on the additive shift constant ($\tau_w\, \alpha_j$). A reform moving τ_w shifts the upper-tail exponent through the additive channel; a credit-supply shift moving the representative LTV would shift it through the gearing channel. The two are observationally separable in cross-LTV-stratified register-data tail-exponent estimation — a sharper identification target than the single-channel reduced-form effects that the existing Norwegian empirical literature reports. Panel (b) reports the cross-type spread $\alpha(z; \text{broken}) - \alpha(z; (\text{C3}))$ across μ-scale factors from 0.5 (low ability) to 2.0 (high ability). The spread is positive at low ability (+0.14 at μ-scale 0.5), crosses zero around μ-scale 1.2, and is negative at high ability (−0.11 at μ-scale 2.0). The sign-change is the substantive new finding: anisotropic Norwegian assessment compresses the bottom of the wealth distribution and concentrates the top relative to the (C3)-uniform counterfactual.

Reading under the post-2022 progressive schedule. The constant rate $\tau_w = 1.1\%$ used in the Pareto-exponent illustration above is exactly the upper-bracket rate of the Norwegian post-2022 two-bracket schedule (Section 4.6). By Proposition 1, the upper-tail exponent $\alpha(z) \approx 2.052$

(unleveraged) and $\alpha(z) \approx 2.041$ (leveraged) reads unchanged when the single-rate calibration is reinterpreted as the two-bracket schedule. Section 5.4's policy invariance follows immediately: changing the lower-bracket rate $\tau_w^{(1)} = 1.0\%$ or the bracket threshold (around NOK 20 M) does not move the upper-tail exponent. On the cross-sectional side, however, the bracket assignment $b(t)$ adds a rank-one mode to the FOC-shift spectrum (Section 4.6): the rank-three structure quoted above (multiplicative-half rank 3 + leveraged additive rank 1) becomes rank-four under the $K = 2$ schedule, with the fourth direction tracing the threshold-crossing households.

The economic reading of the rank addition is a regression-discontinuity- style identification result. Under the post-2022 schedule, the cross-section of post-tax allocations carries an observable bracket-threshold signature: households whose net taxable wealth straddles the NOK 20 M boundary occupy a distinct fourth spectral mode that is absent in any constant-τ_w benchmark. This mode is exploitable as a regression-discontinuity test of the post-tax FOC shift across the threshold, identifying the incremental upper-bracket rate $\tau_w^{(2)} - \tau_w^{(1)}$ from a known regulatory discontinuity. The policy invariance of Proposition 1 adds a second testable prediction: long-run upper-tail inequality measures (top-1% share, top-0.1% share) should respond only to changes in the upper-bracket rate $\tau_w^{(2)} = 1.1\%$, not to threshold shifts or lower-bracket rate changes. This is a falsifiable cross-instrument prediction that distinguishes the Pareto-exponent channel (upper-tail decay rate, governed by $\tau_w^{(K)}$) from the average-tax channel (revenue and middle-of-distribution shape, governed by all bracket parameters).

7 Connection to the corpus

The result of Theorem 1 sits inside a longer chain of results across the corpus, each capturing a different projection of the same FOC shift structure. This section locates the present paper relative to five neighbouring contributions and identifies what each contains, what it depends on, and what the present theorem adds.

7.1 To the cross-sectional spectral framework (Frøseth, 2026h)

Frøseth (2026h) establishes that under (C1)–(C3) with homogeneous returns, the post-tax FOC shift is the rank-one outer product $\mathbf{1}_T \boldsymbol{\delta}^\top$ of Corollary 1, and that the cross-sectional fluctuation spectrum of the allocation matrix is invariant. Theorem 1 relaxes the homogeneous-returns assumption and replaces the rank-one outer product with a row-by-row formula whose total rank is bounded by $\min(T, N) + 1$; the homogeneous case is recovered as Corollary 1. The fluctuation-SDE argument is the same in both papers—(C3) preserves the diffusion $\mathbf{V}$ row-by-row and the fluctuation drift is independent of $\boldsymbol{\mu}_t$—so the spectral content of Frøseth (2026h) carries over to the heterogeneous case for the bulk spectrum away from the rank-$(d+1)$ deterministic shift identified here.

7.2 To the combined flow and stock tax system (Frøseth, 2026b)

The combined operator $\mathcal{T}^{\text{gen}}$ on excess-drift and diffusion variables of Frøseth (2026b) reduces, at the log-wealth level, to the affine pair (λ, c) used throughout this paper. Under (C1)–(C3) with homogeneous returns, Frøseth (2026b) establishes that this pair acts as a uniform rescale-and-shift on all assets simultaneously and preserves the optimal portfolio identity. Theorem 1 extends the row-by-row structure of (λ, c) to heterogeneous returns: the same scalar pair governs both halves of the FOC shift (the multiplicative half through $\lambda - 1$, the additive half through c), but the multiplicative half is no longer a pure rescale of a single direction once $\boldsymbol{\mu}_t$ varies across t.

7.3 To the heterogeneous-returns Fokker–Planck framework (Frøseth, 2026c)

Frøseth (2026c) develops a joint (log-wealth, ability) Fokker–Planck framework for the dynamics under heterogeneous returns and establishes a flow-only differential-incidence result at the distributional level. Theorem 1 gives the deterministic FOC-shift driving those dynamics, and the flow-only sub-system $\tau_w = 0$ of Corollary 3(a) is the cross-sectional companion of Frøseth (2026c)'s incidence statement: the multiplicative-half compression is the spectral signature of the asymmetric flow-tax incidence that Frøseth (2026c) reads off from the joint Fokker–Planck equation. Equilibrium FOC shift here, out-of-equilibrium dynamics there, the same $(\lambda - 1)$ compression in both.

7.4 To the joint Pareto exponent under heterogeneous progressive taxation (Frøseth, 2026d)

Frøseth (2026d) derives the type-conditional joint Pareto exponent in closed form from the demographic-augmented joint Fokker–Planck equation. Theorem 1 gives the FOC-shift ingredient that closes the loop between (C1)–(C3) and the Pareto transformation: under (C1)–(C3) the diffusion $\mathbf{V}$ is preserved row-by-row, the fluctuation drift is invariant (Theorem 1, claim 3), so the entire post-tax / pre-tax contrast in the stationary tail behaviour is mediated by the deterministic FOC shift derived here. The closed-form Pareto transformation of Frøseth (2026d, Theorem 1) is therefore the distributional consequence of the cross-sectional FOC formula (8) of this paper. The bracket-strand recovery argument of Frøseth (2026d, Cor. 2) carries through to Proposition 1 verbatim: the upper-tail Pareto exponent is governed by the topmost bracket alone, with lower brackets shaping only the middle of the distribution.

7.5 To the bracket-strand structure of progressive wealth taxation (Frøseth, 2026f)

Frøseth (2026f) develops the bracket-strand machinery for progressive wealth-tax schedules under (C1)–(C3) with homogeneous returns: piecewise-constant $\tau_w(x)$ generates a piecewise-constant drift on log-wealth, with the steady-state distribution following an asymmetric Laplace form on each bracket and matching the single-bracket Pareto exponent $\alpha_2 = 2(r/k - m_0)/\sigma^2$ in the upper bracket. Corollary 6 extends this machinery in two directions simultaneously. First, it lifts the bracket-strand structure to heterogeneous returns and broken (C3): each bracket contributes a rank-one outer product $\mathbf{1}^{(b)}(\Lambda \mathbf{V}^{-1} \Lambda \mathbf{c}^{(b)})^\top$ with a broken-(C3) shift vector $\mathbf{c}^{(b)}$ and a leverage-conjugated direction. Second, it lifts to the leveraged household balance sheet of Corollary 5: the bracket-conditional structure is *additive* on top of the leverage-conjugated multiplicative half. Proposition 1 then transports the single-bracket Pareto formula of Frøseth (2026f) into the broken-(C3) and leveraged setting through the upper-bracket restriction. The Norwegian post-2022 calibration of Frøseth (2026f, §5) sits as the $K = 2$ instance of this extended structure.

7.6 To the drift-shift symmetry of proportional taxation (Frøseth, 2026g)

Frøseth (2026g) establishes that proportional taxation acts on any Itô diffusion as a pure drift shift, regardless of the return distribution. The stock-only sub-system $\tau_c = \tau_d = 0$ of Corollary 3(b) recovers this symmetry at the allocation-matrix level under heterogeneous returns: the multiplicative-half vanishes ($\lambda = 1$), the additive-half $(c/\gamma)\,\mathbf{1}_T(\mathbf{V}^{-1}\mathbf{1})^\top$ is a row-uniform translation along the global-minimum-variance direction, and cross-agent heterogeneity is preserved entirely. Frøseth (2026g)'s asset-return-level statement and Corollary 3(b) are the same structural fact viewed at two levels: returns transform additively, allocations translate uniformly, the spectral shape of the cross-section is untouched.

8 Discussion

8.1 What the paper resolves and what it leaves open

The paper closes the gap between the homogeneous-returns spectral neutrality of Frøseth (2026b,h) and the actual Norwegian dual income tax across four channels: heterogeneous returns, differential assessment, canonical leverage, and progressive bracket schedules. The two headline theorems and the chain of six corollaries (Sections 3 and 4) carry the deterministic-equilibrium analysis; Theorem 2 and Proposition 1 carry the long-run distributional analysis. The fluctuation-SDE invariance of Theorem 1 (claim 3) preserves the bulk spectrum of the cross-section under all four channels.

Three things are deliberately left aside. First, the spectral content of $\mathbf{V}$ itself: the diffusion is taken as exogenous and shared row-by-row, while in practice ability-stratified portfolio composition could couple ability to $\mathbf{V}$, an interaction not captured here. Second, non-canonical leverage configurations (directly held with binding personal interest cap; the pre-2017 full-deduction sheltering regime), discussed in Section 8.3. Third, the long-horizon spectral object beyond the single-period equilibrium FOC, also discussed in Section 8.3.

8.2 Empirical implications

The Norwegian calibration of Section 6 generates concrete predictions for register-data spectra at each of the four channels.

Heterogeneous returns + (C3). With $T \approx 8$ years, $N \approx 6$ asset classes, and $d \approx 2$ ability factors, the post-tax mean-allocation matrix $\bar{\mathbf{W}}$ should exhibit $\mathrm{rank}(\bar{\mathbf{W}}^{\mathrm{post}} - \bar{\mathbf{W}}) = 3$ at the (C3)-counterfactual benchmark, and deviations from rank three would have to be attributed to either further ability factors or (C3) being violated. The multiplicative-half compression $\|(\bar{\mathbf{W}}^{\mathrm{post}} - \bar{\mathbf{W}})|_{\mathrm{mult}}\| \propto |1-\lambda|$ and the additive-half magnitude $|c|/\gamma$ are separately identified by the spectral decomposition of Corollary 3, allowing the flow- and stock-tax channels to be measured independently rather than as a single reduced-form effect.

Differential assessment. Cor. 4 predicts that under the actual Norwegian asset-class assessment fractions, the additive direction rotates from $\mathbf{V}^{-1}\mathbf{1}$ to $\mathbf{V}^{-1}\mathbf{c}$ but the rank decomposition is unchanged. The empirical signature is therefore the *direction* of the rank-one additive shift, not its rank: a register-data test checking whether the additive component is aligned with $\mathbf{V}^{-1}\boldsymbol{\alpha}_{\text{Norway}}$ versus $\mathbf{V}^{-1}\mathbf{1}$ would distinguish broken (C3) from a counterfactual uniform-assessment system.

Leveraged household balance sheets. Cor. 5 predicts that under the canonical post-2017 leverage configuration, the additive direction further shifts to $\Lambda\mathbf{V}^{-1}\Lambda\mathbf{c}$, with Λ encoding the asset-class LTV pattern. The shift is concentrated in the heavily leveraged asset classes (primary housing); for two households with identical bare-$\boldsymbol{\alpha}$ exposures but different LTVs, Corollary 5 predicts amplification of the equity share in discounted asset classes by the gearing factor while leaving the wealth-tax wedge per unit of equity identical. Tests across LTV strata identify the gearing channel separately from the assessment-fraction channel.

Progressive bracket schedule. Cor. 6 predicts that the post-2022 $K = 2$ schedule adds one rank degree of freedom to the spectrum: rank four (rather than rank three) in the leveraged Norwegian calibration, with the fourth direction tracing the threshold-crossing households. The threshold-discontinuity signature is testable in panel data using households whose taxable wealth straddles the NOK 20 M boundary. Proposition 1 predicts that the upper-tail Pareto exponent depends only on $\tau_w^{(2)}$; this provides a direct identification strategy for the topmost bracket rate from upper-tail wealth-distribution data, independent of the lower bracket and the threshold.

8.3 Future work

The remaining frontier is the long-horizon spectral object. The present paper analyses the cross-sectional spectrum of the allocation matrix at the single-period equilibrium FOC; the horizon-dependent behaviour of the spectrum — and in particular the boundary between the additive (random-matrix-theory) regime and the drift-dominated regime, controlled by the Bouchaud–Potters crossover scale $T_c \sim (\sigma/\mu)^2$ (Bouchaud and Potters, 2003) — is left for future work.

Two configurations of the Norwegian household balance sheet fall outside the canonical case of Section 4.5 and are left for subsequent work. The wealth-tax-side institutional description of both is given in Frøseth (2026a, §4.4, §11); what remains is the cross-sectional and long-run distributional analysis.

The first is the directly held configuration in which the personal interest deduction is capped at the household's gross capital income. When the cap binds, the effective deduction rate falls below τ_c and the (C1) symmetry on the debt margin breaks; the leveraged FOC shift then carries an additional wedge whose magnitude is endogenous to the household's portfolio composition (since gross capital income depends on the dividend-yielding portion of the chosen allocation). The closed form of Corollary 5 does not extend, and a fixed-point characterisation is required.

The second is the pre-2017 full-deduction rule under which leveraged discounted assets shelter taxable wealth from other assets. The wealth-tax base is non-negative at the household level, so a sufficiently leveraged discounted-asset position drives the aggregate base to the zero boundary; the effective wealth-tax rate is then piecewise in the household's post-shelter base, with the boundary determined endogenously by the leveraged composition rather than by an exogenous wealth threshold (as in the bracket schedules of Section 4.6). Theorem 2 requires piecewise patching across the floor, and the eigenvalue ODE becomes regime-conditional.

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.

A Proofs of corollaries

This appendix collects the proofs of Corollaries 1 to 5, deferred from Section 4.

A.1 Proof of Corollary 1

Proof. When $\boldsymbol{\mu}_t \equiv \boldsymbol{\mu}$, the row formula in Theorem 1 collapses to $(\Delta\mathbf{M})_{t\cdot} = (\lambda-1)(\boldsymbol{\mu} - r_f\mathbf{1})^\top + c\,\mathbf{1}^\top$, identical across rows. Stacking gives

$$\Delta\mathbf{M} \;=\; \mathbf{1}_T\left((\lambda-1)(\boldsymbol{\mu} - r_f\mathbf{1}) + c\,\mathbf{1}\right)^\top,$$

a rank-one outer product. Right-multiplication by $\mathbf{V}^{-1}/\gamma$ preserves the rank-one structure and yields (9) with $\boldsymbol{\delta} = (1/\gamma)\,\mathbf{V}^{-1}((\lambda-1)(\boldsymbol{\mu} - r_f\mathbf{1}) + c\,\mathbf{1})$. □

A.2 Proof of Corollary 2

Proof. Substituting (10) into the multiplicative-half $(\lambda-1)(\mathbf{M} - r_f\mathbf{1}_T\mathbf{1}_N^\top)$, column j equals $(\mu_{0,j} - r_f)\,\mathbf{1}_T + \sum_{k=1}^d \nu_{k,j}\,\mathbf{z}_k$, so its column space lies in $\mathrm{span}(\mathbf{1}_T, \mathbf{z}_1, \ldots, \mathbf{z}_d) \subset \mathbb{R}^T$. The joint independence hypothesis makes this span $(d+1)$-dimensional. If $\boldsymbol{\mu}_0 \neq r_f\mathbf{1}$ then at least one column carries a non-zero $\mathbf{1}_T$ component, so $\mathbf{1}_T$ lies in the multiplicative-half's column span and the latter has dimension $\min(d+1, N)$ (capped by the column count N). If $\boldsymbol{\mu}_0 = r_f\mathbf{1}$ then the columns are pure ability combinations, the multiplicative-half's column span equals $\mathrm{span}(\mathbf{z}_1, \ldots, \mathbf{z}_d)$ (a d-dimensional subspace not containing $\mathbf{1}_T$), and its rank is $\min(d, N)$. The additive-half $c\,\mathbf{1}_T\mathbf{1}_N^\top$ has column space $\mathrm{span}(\mathbf{1}_T)$, contained in the multiplicative-half's column span iff $\boldsymbol{\mu}_0 \neq r_f\mathbf{1}$. Adding a rank-one outer-product outside the existing column span increases total rank by one; inside leaves it unchanged. Right-multiplication by $\mathbf{V}^{-1}/\gamma$ preserves rank, yielding the two branches of (11). □

A.3 Proof of Corollary 3

Proof. (a) and (b) by direct substitution of $\tau_w = 0$ and $\tau_c = \tau_d = 0$ into the formulas $\lambda = (1-\tau_c)(1-\tau_d)$ and $c = -(1-\tau_d)\tau_c r_f - \tau_w a$ of (3), applied to the decomposition (8). For (c): the joint Fokker–Planck framework of Frøseth (2026c,d) couples the deterministic equilibrium drift to the type-conditional log-wealth tail behaviour. Under $\mathcal{T}^{\text{gen}}$ the equilibrium drift moves by exactly the per-row shift (7), while diffusion is preserved; substituting this drift change into the type-conditional joint Pareto formula of Frøseth (2026d, Theorem 1) reproduces the closed form they obtain directly from the joint Fokker–Planck equation. □

A.4 Proof of Corollary 4

Proof. Under (C1) and (C2) alone, the after-tax excess return on asset j for agent t retains the form $R^{\text{ex,post}}_{t,j} = \lambda(\mu_{t,j} - r_f) + c_j$ of (4), with $\lambda = (1-\tau_c)(1-\tau_d)$ and c_j given by (13) (the wealth-tax cost on asset j is $\tau_w \alpha_j$ per unit value, asset-specific under broken (C3); the corporate, dividend, and capital-income channels remain asset-uniform under (C1)). Substituting into the per-row mean-variance FOC (Section 2.2) and stacking rows:

$$\bar{\mathbf{W}}^{\text{post}} \;=\; \tfrac{1}{\gamma}\big[\lambda(\mathbf{M} - r_f \mathbf{1}_T \mathbf{1}_N^\top) + \mathbf{1}_T \mathbf{c}^\top\big]\, \mathbf{V}^{-1}.$$

Subtracting the pre-tax FOC $\bar{\mathbf{W}} = (\mathbf{M} - r_f \mathbf{1}_T \mathbf{1}_N^\top)\mathbf{V}^{-1}/\gamma$ yields (12). The rank-decomposition argument of Theorem 1 applies verbatim: the multiplicative half has rank $\operatorname{rank}(\mathbf{M} - r_f \mathbf{1}_T \mathbf{1}_N^\top) \in \{1, \ldots, \min(T,N)\}$, the additive half $\mathbf{1}_T \mathbf{c}^\top/\gamma$ is rank one (an outer product of $\mathbf{1}_T$ with $\mathbf{c}$), and right-multiplication by $\mathbf{V}^{-1}$ preserves rank in each summand. The ability-affine sharpening of Corollary 2 also carries over: the column-span argument used there depends on the structure of $\mathbf{M} - r_f \mathbf{1}_T \mathbf{1}_N^\top$ in the T-direction, which is unchanged, and on the additive half being rank-one along $\mathbf{1}_T$ in the T-direction, which is also unchanged. The (C3)-clean case $\boldsymbol{\alpha} = a\,\mathbf{1}$ recovers $\mathbf{c} = c\,\mathbf{1}$ and Theorem 1 verbatim. □

A.5 Proof of Corollary 5

Proof. Per-agent FOC. Under the canonical case, the post-tax excess return per unit of equity in class j for agent t is $R^{\text{ex,lev}}_{t,j} = \lambda(\mu_{t,j} - r_f)/(1-\ell_j) + c_j$ by (15), with c_j row-uniform under (C1)+(C2). The mean-variance FOC for agent t over equity shares $\boldsymbol{\theta}_t \in \mathbb{R}^N$ uses the equity covariance matrix $\mathbf{V}^{\text{eq}} := \Lambda^{-1}\mathbf{V}\Lambda^{-1}$, giving $\boldsymbol{\theta}^*_t = (\gamma \mathbf{V}^{\text{eq}})^{-1} R^{\text{ex,lev}}_t = \gamma^{-1}\Lambda \mathbf{V}^{-1}\Lambda\, R^{\text{ex,lev}}_t$. Component j of $\Lambda R^{\text{ex,lev}}_t$ is $(1-\ell_j)R^{\text{ex,lev}}_{t,j} = \lambda(\mu_{t,j} - r_f) + (1-\ell_j)c_j$, so $\boldsymbol{\theta}^*_t = \gamma^{-1}\,\Lambda\,\mathbf{V}^{-1}\big[\lambda(\boldsymbol{\mu}_t - r_f \mathbf{1}) + \Lambda \mathbf{c}\big]$.

Stacking rows. Stacking the row vectors $(\boldsymbol{\theta}^*_t)^\top$ as the rows of $\bar{\mathbf{W}}^{\text{post}}$ gives

$$\bar{\mathbf{W}}^{\text{post}} \;=\; \tfrac{\lambda}{\gamma}\,\mathbf{M}_{\text{centered}}\,\mathbf{V}^{-1}\,\Lambda \;+\; \tfrac{1}{\gamma}\,\mathbf{1}_T\,(\Lambda\,\mathbf{V}^{-1}\Lambda\,\mathbf{c})^\top.$$

Subtracting the pre-tax FOC $\bar{\mathbf{W}} = \gamma^{-1}\,\mathbf{M}_{\text{centered}}\mathbf{V}^{-1}$ and collecting the multiplicative term yields (16).

Rank decomposition. The multiplicative half $\mathbf{M}_{\text{centered}}\mathbf{V}^{-1}(\lambda\Lambda - I)/\gamma$ has rank $\operatorname{rank}(\mathbf{M}_{\text{centered}}) \le \min(T,N)$, since right-multiplication by the diagonal $\lambda\Lambda - I$ preserves rank whenever this matrix is non-singular (which holds for $\Lambda \neq \lambda^{-1}I$). The additive half $\mathbf{1}_T(\Lambda\mathbf{V}^{-1}\Lambda\mathbf{c})^\top/\gamma$ is rank one as an outer product. The total rank is therefore at most $\min(T, N) + 1$.

Ability-affine sharpening and unleveraged limit. The ability-affine sharpening of Corollary 2 carries through because the T-direction structure of $\mathbf{M}_{\text{centered}}$ is unchanged and the additive half is rank-one along $\mathbf{1}_T$ in the T-direction. The unleveraged limit $\Lambda = I$ gives $\lambda\Lambda - I = (\lambda - 1)I$ and $\Lambda\mathbf{V}^{-1}\Lambda\mathbf{c} = \mathbf{V}^{-1}\mathbf{c}$, recovering Corollary 4 verbatim. □

A.6 Proof of Corollary 6

Proof. Bracket-conditional FOC. Within bracket b, the wealth-tax rate is $\tau_w^{(b)}$ uniform in wealth, so the post-tax excess return per unit of equity in class j for an agent t with $b(t) = b$ retains the form of (15) with $\mathbf{c}$ replaced by $\mathbf{c}^{(b)}$: $R_{t,j}^{\text{ex,prog}} = \lambda(\mu_{t,j} - r_f)/(1 - \ell_j) + c_j^{(b)}$. The per-agent FOC argument of Corollary 5 then gives $\boldsymbol{\theta}_t^* = \gamma^{-1}\,\Lambda\,\mathbf{V}^{-1}\big[\lambda(\boldsymbol{\mu}_t - r_f\mathbf{1}) + \Lambda\mathbf{c}^{(b(t))}\big]$.

Stacking rows. Stacking the row vectors $(\boldsymbol{\theta}_t^*)^\top$ as the rows of $\bar{\mathbf{W}}^{\text{post}}$, the multiplicative half is bracket-independent (the leverage matrix Λ and the flow-tax retention λ do not depend on b) and stacks to $\frac{\lambda}{\gamma}\,\mathbf{M}_{\text{centered}}\,\mathbf{V}^{-1}\,\Lambda$ exactly as in the constant-τ_w case. The additive half stacks bracket-by-bracket: for rows t with $b(t) = b$, the additive contribution is $\gamma^{-1}(\Lambda\mathbf{V}^{-1}\Lambda\mathbf{c}^{(b)})^\top$, which can be written as the rank-one outer product $\mathbf{1}^{(b)}(\Lambda\mathbf{V}^{-1}\Lambda\mathbf{c}^{(b)})^\top$ since $\mathbf{1}_t^{(b)} = 1$ exactly on the rows where this contribution applies. Summing over brackets and subtracting the pre-tax FOC $\bar{\mathbf{W}} = \gamma^{-1}\,\mathbf{M}_{\text{centered}}\mathbf{V}^{-1}$ yields (18).

Rank decomposition. The multiplicative half has rank $\leq \min(T, N)$ as in the proof of Corollary 5. Each bracket-conditional outer product $\mathbf{1}^{(b)}(\Lambda\mathbf{V}^{-1}\Lambda\mathbf{c}^{(b)})^\top$ is rank one. Summing K such outer products gives an additive half of rank at most K, with equality when the indicator vectors $\{\mathbf{1}^{(b)}\}_{b=1}^K$ are linearly independent in $\mathbb{R}^T$ and the shift directions $\{\Lambda\mathbf{V}^{-1}\Lambda\mathbf{c}^{(b)}\}_{b=1}^K$ are linearly independent in $\mathbb{R}^N$ (which holds whenever the bracket rates are distinct and $\boldsymbol{\alpha} \neq 0$). The total rank of $\bar{\mathbf{W}}^{\text{post}} - \bar{\mathbf{W}}$ is therefore at most $\min(T, N) + K$.

Ability-affine sharpening and reductions. Under the ability-affine hypothesis of Corollary 2 the multiplicative-half rank sharpens to $\min(d+1, N)$, giving total rank $\leq \min(d+1, N) + K$. The $K = 1$ case has a single bracket-indicator $\mathbf{1}^{(1)} = \mathbf{1}_T$ and a single shift $\mathbf{c}^{(1)} = \mathbf{c}$, recovering Corollary 5 verbatim. The $\Lambda = I$ case (no leverage) gives $\lambda\Lambda - I = (\lambda - 1)I$ and $\Lambda\mathbf{V}^{-1}\Lambda\mathbf{c}^{(b)} = \mathbf{V}^{-1}\mathbf{c}^{(b)}$, recovering the bracket-stratified version of Corollary 4. □

References

Maximilien Bernard, Jean-Philippe Bouchaud, and Pierre Le Doussal. A mean-field theory for heterogeneous random growth with redistribution. *arXiv preprint*, 2026. arXiv:2503.23189v3.

Jean-Philippe Bouchaud and Marc Potters. *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management.* Cambridge University Press, 2nd edition, 2003. doi: 10.1017/CBO9780511753893.

James R. Bunch and Christopher P. Nielsen. Updating the singular value decomposition. *Numerische Mathematik*, 31(2):111–129, 1978. doi: 10.1007/BF01397471.

James R. Bunch, Christopher P. Nielsen, and Danny C. Sorensen. Rank-one modification of the symmetric eigenproblem. *Numerische Mathematik*, 31(1):31–48, 1978. doi: 10.1007/BF01396012.

Evsey D. Domar and Richard A. Musgrave. Proportional income taxation and risk-taking. *Quarterly Journal of Economics*, 58(3):388–422, 1944. doi: 10.2307/1882847.

Andreas Fagereng, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. Heterogeneity and persistence in returns to wealth. *Econometrica*, 88(1):115–170, 2020.

Andreas Fagereng, Luigi Guiso, and Marius A. K. Ring. How much and how fast do investors respond to equity premium changes? evidence from wealth taxation. Working paper, 2024.

Anders G. Frøseth. Extensions to the wealth tax neutrality framework. 2026a. arXiv:2603.05277 [physics.soc-ph].

Anders G. Frøseth. Flow taxes, stock taxes, and portfolio choice: A generalised neutrality result. 2026b. arXiv:2603.15974 [physics.soc-ph].

Anders G. Frøseth. Heterogeneous returns and wealth tax neutrality: A Fokker–Planck framework. 2026c. arXiv:2603.16006 [physics.soc-ph].

Anders G. Frøseth. Progressive wealth taxation under heterogeneous returns: A synthesis. Working paper, 2026d.

Anders G. Frøseth. Asset returns, portfolio choice, and proportional wealth taxation. 2026e. arXiv:2603.05264 [physics.soc-ph].

Anders G. Frøseth. Progressive wealth taxation in the combined flow–stock framework: Optimal inequality reduction under gini and top-share targets. Working paper, 2026f.

Anders G. Frøseth. Wealth taxation as a drift modification: A Fokker–Planck approach to tax neutrality. 2026g. arXiv:2603.05283 [physics.soc-ph].

Anders G. Frøseth. Spectral portfolio theory: From SGD weight matrices to wealth dynamics. 2026h. arXiv:2603.09006 [physics.soc-ph].

Agnar Sandmo. Portfolio theory, asset demand and taxation: Comparative statics with many assets. *Review of Economic Studies*, 44(2):369–379, 1977.

Skatteetaten. Verdsettingsrabatt ved fastsetting av formue. `https://www.skatteetaten.no/en/person/taxes/get-the-taxes-right/valuation-discount-in-connection-with-assessment-of-wealth/`, 2026. Accessed February 2026.

G. W. Stewart. Error and perturbation bounds for subspaces associated with certain eigenvalue problems. *SIAM Review*, 15(4):727–764, 1973. doi: 10.1137/1015095.

G. W. Stewart and G. Zhang. Eigenvalues of graded matrices and the condition numbers of a multiple eigenvalue. *Numerische Mathematik*, 58:703–712, 1990. doi: 10.1007/BF01385650.

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