Progressive Wealth Taxation under Heterogeneous Returns: A Synthesis

Anders G. Frøseth*

July 8, 2026

Abstract

Frøseth (2026f, §6.6) report that, under homogeneous returns, a $\sim$ 7-percentage-point dividend-tax increase replicates the structural redistributive role of the Norwegian wealth tax. We test that substitutability claim under heterogeneous returns. Extending the design framework of Frøseth (2026f) to type-stratified $(m_0(z), \sigma(z))$, we (i) prove a generalised neutrality lemma on the heterogeneous $(C1)$–$(C3)$ envelope, (ii) derive a type-conditional joint Pareto formula combining Frøseth (2026c, §5) with the demographic-turnover mechanism of Frøseth (2026f, §6.5), and (iii) introduce formal definitions of an income-flow ratio δ_i and a type-conditional dormant share $\theta(z)$, separating the fraction of wealth flow taxes structurally cannot reach (Domar and Musgrave, 1944; Stiglitz, 1969). A Norwegian recalibration on the Frøseth (2026a) Kapital-400 panel shows that under a common θ the dividend-tax substitute required to absorb the wealth tax rises from $\sim$ 7 to $\sim$ 18 percentage points, and once the fritaksmetoden / participation-exemption channel is taken into account it rises to $\sim$ 40. Below $\theta_H \approx 0.45$ the required compensation exceeds any historical reference, and below a structural threshold $\theta_{\text{struct}} \approx 0.12$ — interior to the fritaksmetoden band — no dividend-tax rate restores the status-quo exponent at all. The substitutability of Frøseth (2026f, §6.6) is therefore conditional, and the condition does not hold in Norway.

JEL Classification: D31, H21, H23, H24, C61.

Keywords: Wealth tax, heterogeneous returns, progressive taxation, Fokker–Planck equation, use-it-or-lose-it, optimal taxation, dormant wealth.

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

Contents

1 Introduction

The Norwegian wealth-tax debate, like its counterparts in several advanced economies, turns on a substitutability question. If the structural redistributive role of the wealth tax could be replicated by a modest adjustment to the flow-tax channel (dividend tax, capital-income tax, corporate tax), the case for retaining the wealth tax becomes narrowly administrative: a question of revenue collection rather than of structural distinctness. If the substitutability fails, the wealth tax has a structural role that flow taxes cannot reach by construction, and the policy choice between the two instruments becomes substantive.

Frøseth (2026f, §6.6) answers the substitutability question affirmatively under the homogeneous-returns assumption: a $\sim$ 7-percentage-point dividend-tax increase suffices to replicate the structural effect of the current Norwegian wealth tax on the upper-tail of the wealth distribution, holding redistributive criteria fixed. The $\sim$ 7-pp magnitude is well within the historical range of Norwegian dividend-tax variation, so the substitutability holds at a politically conceivable cost. The closing paragraph of Frøseth (2026f, §6.6) flags two reasons the result might fail empirically — heterogeneous returns at the top of the distribution, and the share of the wealth-tax base held in dormant assets that generate no current taxable income flow — but does not quantify the failure.

This paper carries out that quantification, finding that the substitutability fails not at the second-order quantitative correction but at the order-of-magnitude level once empirically plausible heterogeneity is admitted. The required flow-tax compensation roughly triples to $\sim$ 18 pp at our Norwegian calibration; under the more accurate *fritaksmetoden*-corrected reading it sits at $\sim$ 40 pp or higher. Neither is within the historical range of Norwegian dividend-tax variation. The wealth tax under realistic Norwegian asset-class composition is therefore not substitutable through flow taxation alone; it has a structurally distinct role.

1.1 The substitutability claim and its caveat

Frøseth (2026f, §6.6) computes a four-row decomposition of the structural redistributive contribution of the current Norwegian wealth tax against the flow-tax channel and the demographic-turnover machinery developed in Frøseth (2026g, §3.2). At calibrated Norwegian parameters, the wealth tax contributes about 17% of the upper-tail compression beyond the no-tax baseline; the flow-tax channel contributes about 83%. Solving for the flow-tax retention factor $k^\dagger$ that delivers the same inequality target (the joint Pareto exponent $\alpha \approx 1.30$ matching the Kapital 400 panel estimate of Frøseth, 2026a, §4) at $r = 0$ gives $k^\dagger \approx 0.43$ versus the status-quo $k = 0.485$, equivalent to a dividend-tax increase from $\sim$ 38% to $\sim$ 45%. The $\sim$ 7-pp adjustment is the substitutability headline: under the homogeneous-returns assumption, the wealth tax's structural effect can be replicated through flow taxes alone.

The caveat noted in the closing paragraph of Frøseth (2026f, §6.6) is that the substitutability calculation depends on two assumptions both of which are violated at the empirical Norwegian top. *First*, the homogeneous-returns assumption embedded in the (C1)–(C3) flow-tax envelope holds neither cross-investor (Fagereng et al., 2020 document persistent return heterogeneity of 5–8 pp at the Norwegian top) nor across asset classes (a substantial share of upper-tail wealth is held in private business holdings under the *fritaksmetoden* participation exemption, where the flow-tax channel is structurally inert). *Second*, the calibration aggregates these two heterogeneity dimensions into a single representative type, which is the form on which the substitutability calculation is well-posed. Disaggregating along either dimension changes the calculation; disaggregating along both produces the order-of-magnitude failure that this paper documents.

1.2 Three classical channels

The structural distinction between flow and stock taxation under heterogeneous returns is classical. Domar and Musgrave (1944) establish that a proportional income tax with full loss offset acts as a silent partner on the *return margin*: the government shares proportionally in gains and losses, and risk-averse investors respond by leveraging up their risky position to restore their pre-tax Sharpe ratio. The partnership operates through the return distribution. Stiglitz (1969) extends this to the wealth tax, which acts as a silent partner on the *position margin*: the government takes a proportional slice of the asset stock itself, while the return distribution per unit of remaining wealth is unchanged. The two partnerships are equivalent under homogeneous returns; under heterogeneous returns they diverge sharply, and the divergence is the structural source of every result in this paper.

Guvenen et al. (2023) formalise the welfare consequence of the divergence. Under heterogeneous productive ability, a revenue-neutral substitution from flow taxation toward wealth taxation reallocates capital from low-ability to high-ability investors — the "use-it-or-lose-it" efficiency channel. Their quantitative calibration to U.S. data delivers an average welfare gain of $\sim 7\%$ of consumption-equivalent for newborns from a revenue-neutral switch from capital-income to wealth taxation, with the optimal wealth-tax exercise returning $\tau_a^{\mathrm{OWT}} \approx 3\%$ (positive, optimal) versus an optimal capital-income tax of $\tau_k^{\mathrm{OKIT}} = -13.6\%$ (a subsidy, optimal). When return heterogeneity is shut down, the optimal capital-income tax flips from a -13.6% subsidy to a $+25.1\%$ positive tax: the case for wealth taxation over capital-income taxation in their framework rests entirely on the heterogeneous-returns channel. Guvenen et al. (2024) develop the theoretical companion — an analytically tractable infinite-horizon model that establishes the use-it-or-lose-it channel, characterises the optimal mix of wealth and capital-income taxes, and adds an innovation/effort channel through which the wealth tax further increases welfare by incentivising high-productivity entrepreneurial activity.

The common thread across the three classical sources is that heterogeneity makes the two tax instruments structurally distinct: under homogeneous returns they are near-substitutes (modulo discounting and risk-bearing considerations), but under heterogeneity they affect different margins. The Fokker–Planck-side analogue of this distinction is the differential-incidence proposition of Frøseth (2026c, §5), which shows that flow taxes compress the cross-type drift gap on log-wealth by the factor $k = (1 - \tau_c)(1 - \tau_d)$, while wealth taxes leave it unchanged (the wealth-tax rate r is type-blind by construction and cancels in the cross-type difference).

1.3 Position in the corpus

The corpus around the present paper consists of four directly relevant predecessors. Frøseth (2026d) establishes the (C1)–(C3) flow-tax decomposition for the homogeneous-returns case. Frøseth (2026g) treats progressive drift design on log-wealth under the Gini criterion, stock-only. Frøseth (2026c) introduces the heterogeneous-returns extension of the neutrality framework, including the differential-incidence proposition that the present paper ports to the design problem. Frøseth (2026f) treats the combined flow + progressive-stock design problem under inequality-relevant criteria (Gini, top share) under the homogeneous-returns assumption, including the substitutability result this paper conditions and reverses.

The present paper combines these four strands. Section 3 establishes the heterogeneous-returns analog of the (C1)–(C3) framework through three preliminary lemmas: type-by-type neutrality on the heterogeneous (C1)–(C3) envelope (Lemma 1), implementability-constrained design when the type z is not directly observable (Lemma 2), and an extension of the envelope to type-conditional asset-class composition that pins down the share of wealth flow taxes structurally cannot reach (Lemma 3). Sections 4–6 then build on this foundation to derive the joint Pareto exponent under heterogeneous returns and the type-conditional decomposition of structural redistributive contributions, and to recalibrate the substitutability calculation against the Frøseth (2026a) Kapital-400 panel.

1.4 Contributions and roadmap

The paper makes the following contributions.

1. A generalised neutrality result on the heterogeneous (C1)–(C3) envelope (Section 3.1): portfolio neutrality of the (C1)–(C3) flow-tax structure survives type-by-type under heterogeneous-returns extension, with type-conditional post-tax SDE (4) on log-wealth.

2. A type-conditional joint Pareto-exponent formula (17) extending Frøseth (2026f, §6.5) to ability-stratified types (Section 4); the population dominant exponent $\alpha^* = \min_z \alpha(z)$ tracks the heaviest-tailed type rather than the population mean (Proposition 1), a differential-incidence statement on the cross-type drift gap (Proposition 2) is the Fokker–Planck-side analogue of Frøseth (2026c, §5)'s result, and a closed-form iso-α^* propagation result (Proposition 3) shows that the cross-type α-gap widens under revenue-neutral substitution toward the wealth tax — the inequality-criterion form of the Guvenen et al. (2023) use-it-or-lose-it channel.

3. Formal definitions of the income-flow ratio δ_i and the type-conditional dormant-wealth share $\theta(z)$ (Definitions 1–3), with operational pin-down for the major Norwegian asset classes including the *fritaksmetoden*-mediated holding-company channel (Bjerksund and Schjelderup, 2021; Bjerksund et al., 2024), formalising the dormant-wealth extension of Lemma 3.

4. A type-conditional decomposition of structural redistributive contributions (Section 5) under both common asset-class composition and the dormant-wealth refinement. The

decomposition makes precise that the cross-type drift gap more than doubles under realistic Norwegian asset-class heterogeneity (factor of 2.4), that the population dominant Pareto exponent α^* shifts from the finite-mean to the infinite-mean regime, and that the wealth-tax marginal contribution to $\alpha(z_H)$ shrinks by half.

5. A type-stratified Norwegian recalibration (Section 6) anchored on the Norwegian register-data evidence of Fagereng et al. (2020), the US comparator Smith et al. (2023), the Kapital 400 panel of Frøseth (2026a, §4), and the NOU 2022:20 (2022) assessment-fraction conventions. The sensitivity table (Table 6) identifies $\theta(z_H) \approx 0.65$ as the threshold where the required flow-tax compensation crosses the politically-conceivable 10 pp range, and $\theta(z_H) \approx 0.45$ as the threshold below which it exceeds 15 pp. Proposition 4 places a further, structural boundary at $\theta_{\text{struct}} \approx 0.12$, below which no dividend-tax rate restores the status-quo exponent. The *fritaksmetoden*-corrected pin-down places empirical Norway at $\theta(z_H) \in [0.10, 0.20]$, well below the first two thresholds and straddling the structural one.

6. The substitutability finding (Sections 6.4–6.5): under empirically plausible Norwegian heterogeneity, the flow-tax compensation required to substitute for the current wealth tax reaches ~ 18 pp on the common-asset-class calibration and ~ 40 pp or higher under the *fritaksmetoden*-corrected pin-down. Both magnitudes are well outside the historical range of Norwegian dividend-tax variation. The substitutability claim of Frøseth (2026f, §6.6) is therefore conditional, and the conditioning on the homogeneous-returns assumption is empirically binding at the Norwegian top of the wealth distribution.

Section 2 introduces the heterogeneous-returns wealth process and the type-mixture FP equation. Section 3 establishes the three preliminary lemmas absorbed into the synthesis. Section 4 derives the type-conditional joint Pareto formula and the population-tail-dominance result. Section 5 carries out the type-conditional decomposition of structural contributions. Section 6 anchors the calibration empirically and identifies the substitutability threshold. Section 7 positions the contribution in the corpus and the external literature. Section 8 closes with the skill-vs-structural-advantage interpretation question and implications for the Norwegian wealth-tax debate.

2 Setup

The heterogeneous-returns model stratifies the population by an ability variable z with distribution $F(z)$ on a measurable space Z. Each type-z investor faces a wealth process whose return primitives are type-conditional: the arithmetic-mean return $\mu(z)$ and the volatility $\sigma(z)$ both depend on z, with the type-conditional log-drift $m_0(z) = \mu(z) - \sigma(z)^2/2$. The wealth-tax schedule $r(\cdot)$ is a function of observable wealth alone (the type z is private), and the flow-tax retention factor $k = (1 - \tau_c)(1 - \tau_d)$ aggregates the corporate and dividend tax rates as in Frøseth (2026d, §2).

Wealth process. On log-wealth $Y = \log X$, the type-z investor's process satisfies the SDE

$$dY_t \;=\; \big[k\, m_0(z) - r(e^{Y_t})\big]\, dt + \sqrt{k}\, \sigma(z)\, dW_t^{(z)}, \tag{1}$$

under the heterogeneous extension of the (C1)–(C3) flow-tax envelope of Frøseth (2026d); the formal derivation is Lemma 1. The dormant-wealth refinement (Lemma 3) replaces k with the type-conditional effective retention $k_{\text{eff}}(z) = 1 - \theta(z)\,(1-k)$, where $\theta(z)$ is the type-conditional income-producing share of the portfolio.

Type-mixture Fokker–Planck equation. Augmenting the type-conditional log-wealth process with a demographic source–sink at type-conditional rate $\delta(z) > 0$ — births at log-wealth y_0, deaths at uniform rate — the joint stationary density $\pi(y, z)$ satisfies the joint Fokker–Planck equation

$$\partial_y\big[(v_{\text{eff}}(z) - r(e^y))\,\pi(y,z)\big] - D_{\text{eff}}(z)\,\partial_y^2 \pi(y,z) + \delta(z)\,\pi(y,z) \;=\; \delta(z)\,F(dz)\cdot \delta_{y_0}(y), \tag{2}$$

where $v_{\text{eff}}(z) = k\,m_0(z) - r$ is the type-conditional drift after flow tax (under a flat-bracket r; the schedule $r(\cdot)$ generalises directly) and $D_{\text{eff}}(z) = k\,\sigma(z)^2/2$ is the type-conditional diffusion. Marginalising in z gives the population log-wealth density $\pi(y) = \int_Z F(dz)\,\pi(y \mid z)$, the object on which the inequality criteria (Gini, top share, population dominant Pareto exponent α^*) are evaluated.

Notation summary. Table 1 collects the symbols used throughout. The conventions follow the cross-paper notation table in Frøseth (2026d, §1) extended to the type-stratified setting; consistency with the rest of the corpus is documented in `references/notation_conventions.md`.

Table 1: Notation used in this paper. The wealth-tax schedule $r(\cdot)$ is a function of observable wealth; the type variable z is private (Assumption 1 below).

Symbol	Meaning	Reference
$z \in Z$	Ability / type variable	Frøseth, 2026c, §5
$F(z)$	Type distribution	
$\mu(z)$, $\sigma(z)$	Type-conditional return primitives	
$m_0(z) = \mu(z) - \sigma(z)^2/2$	Type-conditional log-drift	
$\delta(z)$	Type-conditional demographic turnover rate	Frøseth, 2026a, §4
$\theta(z)$	Type-conditional income-producing share	Definition 3
$k = (1-\tau_c)(1-\tau_d)$	Flow-tax retention factor	Frøseth, 2026d, §2
$k_{\text{eff}}(z) = 1 - \theta(z)(1-k)$	Effective retention with dormant share	(12)
r, $r(\cdot)$	Wealth-tax rate or schedule	Frøseth, 2026f, §2
$\alpha(z)$	Type-conditional Pareto exponent	(17)
$\alpha^* = \min_z \alpha(z)$	Population dominant exponent	Proposition 1

3 Three preliminary results absorbed as lemmas

This section establishes the heterogeneous-returns analog of the (C1)–(C3) framework through three lemmas. The first extends the proportional-tax neutrality of Frøseth (2026d, Theorem 1) to type-stratified return primitives (Lemma 1), giving the type-conditional post-tax SDE on log-wealth that all subsequent results inherit. The second poses the design problem when the type variable z is unobservable to the tax authority, and shows that under the implementability constraint the marginal-density extremisation reduces to a fixed-point problem in the schedule (Lemma 2). The third extends the framework to type-conditional asset-class composition by introducing the income-flow ratio δ_i and the type-conditional dormant share $\theta(z)$, formalising the channel by which flow taxes structurally cannot reach the dormant-wealth base (Lemma 3). All three lemmas are proved in Appendix A.

3.1 Generalised neutrality on the heterogeneous $(C1)$–$(C3)$ envelope

We combine the flow-tax decomposition of Frøseth (2026d) with the type-mixture treatment of Frøseth (2026c) into a single heterogeneous-neutrality statement, then read off the post-tax SDE on log-wealth that flows through to the type-conditional joint Pareto formula in Section 4 and the decomposition in Section 5.

3.1.1 The heterogeneous (C1)–(C3) class

Each investor is characterised by an ability state $z \in Z$ with population distribution $F(z)$, and faces a portfolio of risky assets with type-conditional gross return $\mu(z)$ and volatility $\sigma(z)$, plus a risk-free asset with rate r_f. The same flow-tax structure (corporate τ_c, capital- income τ_k, dividend τ_d, with assessment fractions α_i for the wealth tax) is assessed type-blindly. The three structural conditions of Frøseth (2026d, §4) carry over verbatim:

(C1) $\tau_k = \tau_c$ (capital-income equals corporate);

(C2) $r_s = r_f$ (shielding rate equals risk-free);

(C3) $\alpha_i = \alpha$ for all assets (uniform assessment).

None of (C1)–(C3) refers to the type variable z; the conditions act on the tax-system structure, not on the investor heterogeneity.

3.1.2 Heterogeneous neutrality

Lemma 1 (Heterogeneous generalised neutrality)**.** *Suppose the tax system satisfies* (C1)–(C3) *as above, and that each type $z \in Z$ optimises a CRRA expected-utility problem over the same tax-blind portfolio menu. Let $\boldsymbol{w}^*(z)$ denote the optimal portfolio weights for type z. Then $\boldsymbol{w}^*(z)$ is independent of all tax rates (τ_c, τ_d, τ_w) for every z.*

Proof. See Appendix A.1. □

The consequence we need downstream is the post-tax wealth process. Under (C1)–(C3), the type-z investor's wealth satisfies

$$\frac{\mathrm{d}X_t}{X_t} = [k\,\mu(z) - r(X_t)]\,\mathrm{d}t + \sqrt{k}\,\sigma(z)\,\mathrm{d}W_t^{(z)}, \tag{3}$$

where $W^{(z)}$ is a type-conditional standard Brownian motion and $r(\cdot)$ is the type-blind progressive wealth-tax schedule. On log-wealth $Y = \log X$, by Itô,

$$\mathrm{d}Y_t \;=\; \left[k\, m_0(z) - r(e^{Y_t})\right] \mathrm{d}t + \sqrt{k}\,\sigma(z)\,\mathrm{d}W_t^{(z)}, \tag{4}$$

with $m_0(z) = \mu(z) - \sigma(z)^2/2$ the type-conditional geometric-mean drift on log-wealth.

Remark (Convention note). Equations (3)–(4) follow the post-tax SDE form of Frøseth (2026f, eq. post_tax_X), in which the (C1)–(C3) flow-tax retention k rescales drift by k and diffusion by k (so the volatility on log-wealth becomes $\sqrt{k}\,\sigma$). This is the convention adopted throughout Frøseth (2026f) and inherited here. An alternative convention used in Frøseth (2026c, eq. combined_z) keeps the diffusion at $\sigma(z)$ without the $\sqrt{k}$ factor, reflecting a different reading of the (C1)–(C3) post-tax SDE. The two conventions differ on the diffusion coefficient but agree on the drift; the qualitative results of Sections 4–5 are unchanged under either reading, with quantitative calibration values shifting modestly between conventions. We keep the Frøseth (2026f) convention for cross-paper consistency with the homogeneous-baseline calculation.

Remark (Cross-type asymmetry of incidence). The heterogeneity-blind structure of (C1)–(C3) means that the flow-tax retention k and the wealth-tax rate r apply type-independently. Their differential effects on cross-type inequality therefore arise entirely through the type-conditional return primitives $(\mu(z), \sigma(z))$ entering (4), and through the type-conditional asset-class composition formalised in Section 3.3 below. The sub-class on which the differential incidence is sharpest is the one where the cross-type spread of $\mu(z)$ is large; this is the channel formalised in Frøseth (2026c, §5) as the differential-incidence proposition (where the flow-tax retention is denoted λ; we identify $\lambda \equiv k$ throughout), recovered in Section 4 as the cross-type drift gap of equation (eq. joint_pareto_het).

3.2 Heterogeneous stock-only design under implementability

Lemma 1 delivers the type-conditional post-tax SDE (4); the next step is to characterise the population steady state and the design problem on it. Under heterogeneity, the planner faces a structural constraint absent from the homogeneous case: the type variable z is private information. The schedule $r(\cdot)$ must be measurable in observable wealth alone, not in the (z, x) pair. The design problem is therefore a constrained optimum within the homogeneous-class machinery of Frøseth (2026g,d,f), with the constraint binding through a type-mixing channel that we make explicit here.

3.2.1 The type-mixture FP equation

Augment (4) with the demographic source-sink at type-conditional rate $\delta(z)$, allowing the turnover rate to vary with type (the homogeneous case sets $\delta(z) \equiv \delta$):

$$\partial_t \pi(y, z, t) \;=\; -\left[k\, m_0(z) - r(e^y)\right] \partial_y \pi + \tfrac{1}{2}\, k\, \sigma(z)^2\, \partial_y^2 \pi - \delta(z)\, \pi + \delta(z)\, \phi(y, z), \tag{5}$$

where $\phi(y,z)$ is the type-preserving reset distribution at the population mean of type z. The marginal density on log-wealth, integrated over types, is

$$\pi(y,t) \;=\; \int_Z F(\mathrm{d}z)\,\pi(y,z,t), \tag{6}$$

and is the object on which redistributive criteria (Gini, top share) are evaluated.

3.2.2 Implementability and the wealth-conditional type distribution

The natural design problem is: choose $r:(0,\infty)\to[0,r_{\max}]$, measurable in raw wealth alone, to extremise the inequality criterion on the marginal π subject to revenue and feasibility constraints. The implementability constraint is

Assumption 1 (Wealth-measurable schedule). The wealth-tax schedule r is measurable with respect to the σ-algebra generated by raw wealth X alone, not the (z,X) pair: the planner cannot condition rates on the type z.

To see how this constraint binds, integrate (5) in z against $F(\mathrm{d}z)$:

$$\partial_t\pi(y,t) \;=\; -\partial_y\big[\bar v(y,t)\,\pi(y,t)\big] + \tfrac{1}{2}\,\partial_y^2\big[\bar D(y,t)\,\pi(y,t)\big] - \bar\delta(y,t)\,\pi(y,t) + \bar\phi(y,t), \tag{7}$$

where $\bar v$, $\bar D$, $\bar\delta$, $\bar\phi$ are the type averages *conditional on observed log-wealth*:

$$\bar v(y,t) = k\,\mathbb{E}[m_0(z)\mid Y_t=y] - r(e^y), \quad \bar D(y,t) = \tfrac{1}{2}\,k\,\mathbb{E}[\sigma(z)^2\mid Y_t=y], \tag{8}$$

with $\bar\delta(y,t)=\mathbb{E}[\delta(z)\mid Y_t=y]$ and $\bar\phi(y,t)=\mathbb{E}[\delta(z)\,\phi(y,z)\mid Y_t=y]/\bar\delta(y,t)$. The conditional type distribution $F(z\mid Y_t=y)$ is the central new object: it is the inverse-image distribution that enters every type-averaged moment. In steady state it satisfies a Bayes-style relation,

$$F(z\mid Y=y) \;=\; \frac{F(z)\,\pi(y\mid z)}{\pi(y)}, \tag{9}$$

which is implicit in r through the steady-state densities $\pi(y\mid z)$ and $\pi(y)$.

3.2.3 The constrained design problem as a fixed point

Lemma 2 (Implementability-constrained design)**.** *Under Assumption 1 and the heterogeneous* $(C1)$–$(C3)$ *envelope of Section 3.1, the inequality-criterion design problem reduces to the fixed-point problem: find* $r:(0,\infty)\to[0,r_{\max}]$ *such that the marginal log-wealth density* $\pi(y;r)$ *from* (7) *extremises a redistributive criterion* $\Phi[\pi]$ *subject to a revenue constraint, where the type-conditioning moments* $\mathbb{E}[m_0(z)\mid Y=y]$ *and* $\mathbb{E}[\sigma(z)^2\mid Y=y]$ *are themselves functionals of* r *via the wealth-conditional type distribution* (9)*.*

Proof. See Appendix A.2. ☐

Remark (Cost of implementability). The unconstrained type-aware design problem permits $r(z,x)$, schedules indexed jointly on type and wealth. Such schedules are not implementable when z is private. The constrained schedule $r(x)$ admits a degenerate decomposition $r(x) = \mathbb{E}_F[r^*(z,x)]$, but this type-blind average is generally *not* the constrained-optimal schedule; the correct constrained optimum accounts for $F(z \mid Y = y)$ via (8). The gap between the unconstrained and constrained optima is the FP-side cost of implementability. We do not characterise it analytically here. The Guvenen et al. (2023) use-it-or-lose-it reading of this gap is that the constrained schedule under-rewards high-ability types relative to the unconstrained schedule.

Remark (Recovery of the homogeneous case). At $z \equiv z_0$ (degenerate type distribution), (8) collapses to $\bar{v}(y) = k\, m_0(z_0) - r(e^y)$ and $\bar{D}(y) = k\, \sigma(z_0)^2/2$, both independent of y. The fixed-point structure of Lemma 2 becomes trivial, and the design problem reduces to the Frøseth (2026f) design problem in Section 4 with the homogeneous parameters (m_0, σ) replaced by $(m_0(z_0), \sigma(z_0))$. The fixed-point machinery is the precise content that heterogeneity adds.

Remark (Wealth-conditional type distribution at the upper tail). For the upper tail of the wealth distribution, the wealth-conditional type distribution $F(z \mid Y = y)$ concentrates on the heaviest-tail type z^* identified in Section 4 as $y \to \infty$, because the heaviest tail dominates the marginal density there. In particular, $\mathbb{E}[m_0(z) \mid Y = y] \to m_0(z^*)$ and $\mathbb{E}[\sigma(z)^2 \mid Y = y] \to \sigma(z^*)^2$ as $y \to \infty$. The upper-tail Pareto exponent of the marginal π is therefore the type-z^* exponent $\alpha(z^*)$, recovering equation (7) of the type-conditional derivation (`notes/derivation_type_conditional_pareto.md` §5).

3.3 Dormant-wealth extension

The two preceding lemmas treat heterogeneity in the *return primitives* $(\mu(z), \sigma(z))$ but assume that all wealth, regardless of asset class, generates taxable income at the rate the (C1)–(C3) decomposition presupposes. In practice the upper tail of the wealth distribution holds a substantial share in assets that produce little or no current-period taxable income: closely-held businesses (retained earnings rather than dividends), real estate (untaxed imputed rent and capital gains), illiquid private holdings, art and other collectibles. Such *dormant wealth* is the share of the upper-tail tax base that flow taxes structurally cannot reach. The wealth tax, assessed on stock value, reaches it uniformly. This is the structural distinction at the heart of Domar and Musgrave (1944); Stiglitz (1969); we make it precise on the (C1)–(C3) envelope here.

3.3.1 Income-flow ratio of an asset class

Let an investor's portfolio be decomposed into asset classes indexed by $i \in \{1, \ldots, K\}$, with wealth share $w_i \in [0,1]$ in class i ($\sum_i w_i = 1$) and gross expected return μ_i on class i.

Definition 1 (Income-flow ratio)**.** For asset class i, the *income-flow ratio* $\delta_i \in [0,1]$ is the fraction of the gross return μ_i realised as currently taxable income flow per unit time — the fraction subject to the corporate, dividend, or capital-income channel of the (C1)–(C3) decomposition. The complementary share $1 - \delta_i$ is the *dormant fraction*: the part of the return that accrues as unrealised appreciation, retained corporate earnings, deferred capital gains, or imputed consumption value, and is not subject to the flow-tax retention factor $k = (1 - \tau_c)(1 - \tau_d)$.

The income-flow ratio is an institutional parameter fixed by the tax statute, not a behavioural choice of the investor — *conditional on the ownership structure*. The same underlying asset can have very different δ_i depending on how it is held, because the participation-exemption regimes common in many advanced economies (the Norwegian *fritaksmetoden*, the EU Parent–Subsidiary Directive, similar regimes in Switzerland, the Netherlands, the UK, Singapore) make corporate-level dividends and capital gains tax-exempt when received by another corporate entity. Listed equity held *directly* therefore has a substantially different δ_i than the same listed equity held *through an unlisted holding company* under the exemption regime. Table 2 reports the operational pin-down for the dominant Norwegian institutional configurations (NOU 2022:20, 2022; Bjerksund and Schjelderup, 2021; Bjerksund et al., 2024; cf. Frøseth 2026a, §6.2, Frøseth 2026e, §9).

Table 2: Operational pin-down of the income-flow ratio δ_i for the major asset classes in the Norwegian wealth-tax base, conditional on ownership structure. Listed equity is split into direct holdings, pension-wrapper holdings, and holdings through unlisted holding companies under the *fritaksmetoden* (participation exemption); the third configuration is the dominant pattern at the top of the Norwegian wealth distribution. The demographic- turnover rate $\delta(z)$ of Section 4 captures the realisation-event flow at liquidation, inheritance, or sale separately and is not double-counted in δ_i.

Asset class and ownership structure	δ_i	Reason
Bank deposits	≈ 1	Interest fully taxed annually
Bonds	≈ 1	Coupon income fully taxed
Listed equity, direct taxable account	≈ 0.4–0.6	Dividends and realised gains taxed; retained earnings deferred
Listed equity, pension wrapper	≈ 0	All returns deferred to withdrawal
Listed equity, via holding co. (*fritaksmetoden*)	≈ 0	Corporate-level dividends and capital gains exempt; personal tax only on distribution to individual, deferrable indefinitely
Closely-held operating business	≈ 0–0.2	Dominantly retained earnings; minimal current dividend payout
Rental real estate	≈ 0.3–0.5	Rental income taxed; capital gains often deferred
Primary residence	≈ 0	Imputed rent untaxed; capital gains often exempt
Art and collectibles	≈ 0	No current income; gains taxed only on realisation

Remark (The participation-exemption channel). The third row of Table 2 is quantitatively the most consequential for the Norwegian top of the wealth distribution. Most substantial Norwegian investors hold their listed-equity exposure through unlisted holding companies (*aksjeselskap*, AS) rather than directly (Frøseth, 2026e, §9.2, Bjerksund and Schjelderup, 2021). Within the holding company, dividends received and capital gains realised on qualifying shareholdings are exempt from corporate tax under the *fritaksmetoden*; the personal capital-income tax (at the $\sim 38\%$ effective rate that includes the shareholder-model uplift) is triggered only when the owner distributes the proceeds out of the holding company. Holding companies therefore accumulate large pools of retained earnings and unrealised gains with personal-tax liability deferred indefinitely

(Frøseth, 2026a, §5). The institutional consequence is that the operative δ_i for listed-equity holdings of substantial Norwegian investors is closer to 0 than to the 0.4–0.6 of direct-holding row two; the listed-equity portfolio behaves in the FP framework like a near-dormant asset rather than an income-producing one. Similar structures exist in most other jurisdictions with participation-exemption regimes (the EU Parent–Subsidiary Directive harmonises this across EU member states; Switzerland and several non-EU jurisdictions have analogues), so the channel is not Norway-specific.

Remark (Implication for the empirical $\theta(z_H)$ pin-down). The Section 3.3 and Section 5.1 calibration $\theta(z_H) \approx 0.4$ uses the direct-holding δ_i for listed equity. Updating to the *fritaksmetoden*-mediated $\delta_i \approx 0$ for the dominant top-of-distribution holding pattern pushes $\theta(z_H)$ substantially lower — plausibly to $\theta(z_H) \in [0.1, 0.2]$, depending on the share of direct-holding exposure retained alongside the holding-company structure. Section 6.3 reports the sensitivity table over the full range $\theta(z_H) \in [0.2, 1.0]$; the infeasibility finding ($\Delta\tau_d > 15$ pp at the common-θ pin-down $\theta_H = 0.4$, well outside the historical Norwegian range) is therefore conservative relative to the more accurate *fritaksmetoden*-corrected calibration, which would place $\theta(z_H) \approx 0.15$ in the $\Delta\tau_d \approx +30$ pp region of Table 6. Even at the most favourable end of the empirical range, the substitutability of Frøseth (2026f, §6.6) fails by an order of magnitude.

3.3.2 Type-conditional dormant-wealth share

Definition 2 (Portfolio-weighted income-producing share)**.** For an investor with portfolio weights $\{w_i\}$ and asset-class income-flow ratios $\{\delta_i\}$, the *portfolio-weighted income-producing share* is

$$\theta \;=\; \sum_{i=1}^{K} w_i \, \delta_i \;\in\; [0,1], \tag{10}$$

and the *portfolio-weighted dormant-wealth share* is $1 - \theta$.

Definition 3 (Type-conditional dormant share)**.** Under heterogeneous returns indexed by ability z with type-conditional portfolio composition $\{w_i(z)\}$, the *type-conditional income-producing share* is

$$\theta(z) \;=\; \sum_{i=1}^{K} w_i(z) \, \delta_i. \tag{11}$$

The wealth tax on raw wealth at rate r acts type-blindly (by (C3)); the flow-tax channel, by contrast, reaches only the type-conditional income-producing share $\theta(z)$ of each type's portfolio.

The framework's heterogeneity loads entirely on the portfolio-share heterogeneity $\{w_i(z)\}$: $\theta(z)$ varies with z because portfolio composition varies with ability, not because the institutional parameter δ_i varies with ability.

3.3.3 Effective retention under asset-class heterogeneity

Aggregating the income-producing and dormant shares of a type-z investor's portfolio, the after-flow-tax expected return on raw wealth is $\theta(z) \cdot k\,\mu(z) + (1-\theta(z)) \cdot \mu(z) = \mu(z) \cdot [1 - \theta(z)(1-k)]$.

Define the type-conditional *effective retention factor*

$$k_{\text{eff}}(z) \;=\; 1-\theta(z)\,(1-k) \;\in\; [k,\,1], \tag{12}$$

which interpolates between the fully-income-producing case $k_{\text{eff}} = k$ at $\theta(z) = 1$ and the fully-dormant case $k_{\text{eff}} = 1$ at $\theta(z) = 0$.

Lemma 3 (Heterogeneous neutrality with dormant wealth)**.** *Suppose Lemma 1 holds for the income-producing share of each type's portfolio, and that the wealth tax assesses raw wealth uniformly across asset classes (so (C3) extends to dormant assets). Let $\theta(z)$ be the type-conditional income-producing share. Then the post-tax wealth dynamics for type z are*

$$\mathrm{d}Y_t \;=\; [k_{\text{eff}}(z)\,m_0(z) - r(e^{Y_t})]\,\mathrm{d}t + \sqrt{k_{\text{eff}}(z)}\,\sigma(z)\,\mathrm{d}W_t^{(z)}, \tag{13}$$

where $k_{\text{eff}}(z)$ is given by (12).

Proof. See Appendix A.3. □

Remark (The dormant-wealth channel widens the drift gap)*.* The cross-type drift gap on log-wealth, conditional on a single bracket r, is now

$$v_{\text{eff}}(z_H) - v_{\text{eff}}(z_L) = k_{\text{eff}}(z_H)\,m_0(z_H) - k_{\text{eff}}(z_L)\,m_0(z_L), \tag{14}$$

which generalises the homogeneous-θ gap $k\,[m_0(z_H)-m_0(z_L)]$ implicit in Lemma 1. The wealth-tax rate r still cancels (the (C3) extension is type-blind in the wealth tax). When $\theta(z_H) < \theta(z_L)$ — the case where high-ability types hold more dormant wealth — $k_{\text{eff}}(z_H) > k_{\text{eff}}(z_L)$, and the flow-tax compression of the gap is *weaker* for the high-ability type than for the low-ability type. The drift gap (14) is therefore wider than the homogeneous-θ formula $k\,[m_0(z_H)-m_0(z_L)]$ predicts: dormant wealth at the top shields high-ability types from the flow-tax compression that the homogeneous-returns calibration of Frøseth (2026f) assumes.

Remark (Norwegian empirical magnitudes)*.* The Kapital 400 panel evidence in Frøseth (2026a, §4) documents that closely-held business holdings, real estate, and illiquid private positions dominate the top-of-distribution wealth composition. A first-order pinning of $\theta(z)$ at the upper tail (the $z = z_H$ regime) sits below 0.5; for the lower tail (z_L, dominated by listed equity, deposits, and primary real-estate residence under the partial Norwegian assessment discount), $\theta(z_L)$ sits closer to 0.7–0.9. Plugging in $k = 0.485$ (Norwegian status quo) gives effective retentions $k_{\text{eff}}(z_L) \approx$ 0.6–0.75 and $k_{\text{eff}}(z_H) \gtrsim 0.75$. The dormant-wealth gap $k_{\text{eff}}(z_H) - k_{\text{eff}}(z_L)$ is on the order of 0.05–0.15, modest in absolute terms but on the same order as the cross-type drift correction $k\,[m_0(z_H) - m_0(z_L)]$ at empirical Kapital 400 heterogeneity. A type-stratified empirical pin-down is the calibration target of Section 6.

Remark (Recovery and limits)*.* Setting $\theta(z) \equiv 1$ for all z recovers Lemma 1's post-tax SDE (4) verbatim, and (12) reduces to $k_{\text{eff}} = k$. Setting $\theta(z) \equiv 0$ for all z recovers the no-flow-tax dynamics $\mathrm{d}Y = (m_0(z) - r)\mathrm{d}t + \sigma(z)\mathrm{d}W$, with $k_{\text{eff}} = 1$. The lemma interpolates linearly in $\theta(z)$

between the two limits.

4 Joint Pareto exponent under heterogeneous returns

The three preliminary lemmas of Section 3 place the heterogeneous design problem on the same analytical footing as the homogeneous case: a type-conditional post-tax SDE (4) (or its dormant-wealth refinement (13)), governed by an implementability constraint that makes the design problem a fixed point in r through the wealth-conditional type distribution (9). We now characterise the upper tail of the type-conditional and population steady states.

4.1 The type-conditional joint Pareto formula

Following Gabaix (1999); Gabaix et al. (2016), augment (4) with the demographic source-sink at type-conditional rate $\delta(z) > 0$. On the upper bracket where $r(x) = r$ is constant, the type-conditional FP equation reduces to the eigenvalue ODE used in Frøseth (2026f, §6.5), with type-conditional coefficients:

$$D_{\text{eff}}(z)\,\pi''(y \mid z) - v_{\text{eff}}(z)\,\pi'(y \mid z) - \delta(z)\,\pi(y \mid z) \;=\; 0, \tag{15}$$

with effective drift and diffusion

$$v_{\text{eff}}(z) = k\,m_0(z) - r, \qquad D_{\text{eff}}(z) = \tfrac{1}{2}\,k\,\sigma(z)^2. \tag{16}$$

Theorem 1 (Type-conditional joint Pareto exponent)**.** *Under Lemma 1 and a type-conditional demographic source-sink at rate $\delta(z) > 0$, each type $z \in Z$'s upper-tail log-wealth marginal $\pi(y \mid z)$ decays as $\pi(y \mid z) \sim A(z)\,e^{-\alpha(z)\,y}$ for y large, with type-conditional Pareto exponent*

$$\alpha(z) \;=\; \frac{-\,v_{\text{eff}}(z) + \sqrt{v_{\text{eff}}(z)^2 + 4\,D_{\text{eff}}(z)\,\delta(z)}}{2\,D_{\text{eff}}(z)}, \tag{17}$$

where $v_{\text{eff}}(z)$ and $D_{\text{eff}}(z)$ are given by (16).

Proof. See Appendix A.4. □

Three consistency checks pin (17) down relative to the homogeneous baseline and the bracket strand of Frøseth (2026f).

Corollary 1 (Recovery of the homogeneous formula)**.** *At a single representative type $z \equiv z_0$, formula (17) reduces to the homogeneous joint Pareto formula of Frøseth (2026f, §6.5) with parameters $(m_0(z_0), \sigma(z_0))$ replacing (m_0, σ). At calibrated Norwegian parameters ($\mu = 0.0992$, $\sigma = 0.28$, $k = 0.485$, $r = 0.006$, $\delta = 0.062$) formula (17) returns $\alpha(z_0) = 1.298$, matching the Frøseth (2026f, §6.6) status-quo value of* 1.30.

Corollary 2 (Recovery of the bracket-strand exponent)**.** *In the no-turnover limit $\delta(z) \to 0$ with finite-mean condition $r/k > m_0(z)$, formula (17) reduces to the bracket-strand Pareto exponent*

$\alpha_2(z) = 2(r/k - m_0(z))/\sigma(z)^2$ *of* Frøseth *(*2026f*, Definition 3) evaluated at the type-conditional regime parameters.*

Proof. See Appendix A.6. □

Corollary 3 (Dormant-wealth refinement)**.** *Under Lemma 3 (asset-class heterogeneity in $\theta(z)$), formula* (17) *extends verbatim with k replaced by the type-conditional effective retention $k_{\text{eff}}(z) = 1 - \theta(z)\,(1-k)$ of* (12)*. The flow-tax channel enters $\alpha(z)$ through $k_{\text{eff}}(z)$; the wealth-tax channel enters through r; the demographic channel enters through $\delta(z)$.*

4.2 Population upper-tail behaviour

The marginal log-wealth distribution on the population is the type average

$$\pi(y) \;=\; \int_Z F(\mathrm{d}z)\,\pi(y \mid z), \tag{18}$$

which is the object the inequality criteria (Gini, top share) of Section 2 are evaluated on. For y large the integrand is dominated by the type whose exponent $\alpha(z)$ is the smallest, since the heaviest-tailed type has the most mass at high log-wealth. We formalise this as follows.

Proposition 1 (Population dominant Pareto exponent)**.** *Suppose the type space Z is finite or compact and the type-conditional Pareto exponent $z \mapsto \alpha(z)$ of* (17) *attains its infimum at some $z^* \in \operatorname{supp}(F)$. Then the population marginal log-wealth distribution* (18) *satisfies $\pi(y) \sim A^* e^{-\alpha^* y}$ as $y \to \infty$, with $\alpha^* = \alpha(z^*)$ and $A^* = F(\{z^*\}) \cdot A(z^*)$ when z^* is an atom of F (or the limit of a continuous density at the inf-attaining point).*

Proof. See Appendix A.7. □

Remark (The use-it-or-lose-it tail dominates)*.* At common volatility $\sigma(z) \equiv \sigma$, the type-conditional exponent (17) is strictly decreasing in $\mu(z)$: higher $\mu(z)$ gives a positive shift in $v_{\text{eff}}(z)$, which decreases $\alpha(z)$ via the $-v_{\text{eff}} + \sqrt{v_{\text{eff}}^2 + \cdots}$ structure. The type with the highest μ has the heaviest tail. The population upper tail is therefore dominated by the high-ability subpopulation that Guvenen et al. (2023) identify as the source of the use-it-or-lose-it efficiency channel; the inequality-criterion analogue is that the high-ability tail is the binding object for top-share calculations.

Remark (Phase transition at population level)*.* α^* may sit below or above 1 even when the homogeneous-baseline $\alpha(z_0) > 1$. At a representative two-type calibration with $\mu(z_L) = 0.07$, $\mu(z_H) = 0.16$, common $\sigma = 0.28$, $k = 0.485$, $r = 0.006$, $\delta = 0.062$, formula (17) gives $\alpha(z_L) = 1.586$ and $\alpha(z_H) = 0.892$, so $\alpha^* = 0.892 < 1$ — the population is in the infinite-mean Pareto regime. At the same calibration the homogeneous baseline gives $\alpha = 1.30 > 1$. The empirical estimate $\hat{\alpha} \approx 1.3$ from the Kapital 400 panel of Frøseth (2026a, §4) is therefore an *average over a heterogeneous population* whose upper tail is plausibly an infinite-mean Pareto. This is the recalibration target of Section 6.

The same phase structure arises in the mean-field model of Bernard et al. (2026), in which agents carry quenched heterogeneous growth rates and interact through a mean-field redistribution term: at a critical redistribution intensity the stationary distribution passes between a localised (condensed) phase, concentrated on the fastest growers, and a delocalised phase, with the tail exponent in the partially localised phase depending on the *variance* of the growth-rate distribution. The ability-stratified formula (17) is the Fokker–Planck counterpart of that mean-field picture: the infinite-mean regime $\alpha^* < 1$ echoes their condensed phase, with demographic turnover δ and the tax channels playing the role that the redistribution intensity plays there.

4.3 The cross-type drift gap

The cross-type drift gap on log-wealth, conditional on a single bracket r, encodes the differential-incidence channel of Frøseth (2026c, §5) in the heterogeneous (C1)–(C3) framework:

$$\Delta v(z_H, z_L) \;=\; v_{\text{eff}}(z_H) - v_{\text{eff}}(z_L) \;=\; k\left[m_0(z_H) - m_0(z_L)\right], \tag{19}$$

which is independent of r (the wealth-tax rate cancels because r acts type-blindly) and scales linearly with k (the flow tax compresses the gap proportionally to its retention). At common σ, $\Delta v = k\,[\mu(z_H) - \mu(z_L)]$. Under the dormant-wealth refinement of Corollary 3 the gap generalises to (14), and the wealth-tax neutrality in r is preserved, but the flow-tax compression becomes type-conditional through $k_{\text{eff}}(z_H)$ versus $k_{\text{eff}}(z_L)$.

Proposition 2 (Differential incidence in the cross-type drift gap)**.** *Under common volatility $\sigma(z) \equiv \sigma$ and common turnover rate $\delta(z) \equiv \delta$, the cross-type log-wealth drift gap $\Delta v(z_H, z_L)$ of (19) responds to the tax instruments as*

$$\frac{\partial \Delta v}{\partial r} = 0, \qquad \frac{\partial \Delta v}{\partial k} = m_0(z_H) - m_0(z_L) > 0. \tag{20}$$

The wealth-tax rate enters type-blindly (a position-margin partner in the sense of Stiglitz 1969); the flow-tax retention factor enters type-dependently through $m_0(z)$ (a return-margin partner in the sense of Domar and Musgrave 1944). The drift gap is invariant to the wealth-tax level and scales linearly with the flow-tax retention factor.

The propagation from this v-level statement to the inequality-criterion level — the cross-type Pareto-exponent gap $\Delta\alpha := \alpha(z_L) - \alpha(z_H) > 0$ — is non-trivial because α is a non-linear function of v through the joint formula (17). The non-linearity changes the policy-relevant signs and is the content of the next proposition.

Proposition 3 (α-gap response to the tax instruments)**.** *Let $S(z) := \sqrt{v_{\text{eff}}(z)^2 + 4\,D_{\text{eff}}\,\delta} = 2\,D_{\text{eff}}\,\alpha(z) + v_{\text{eff}}(z) > 0$. Under common σ and δ and assuming $\mu(z_H) > \mu(z_L)$:*

1. ***Unconstrained partial derivatives.***

$$\frac{\partial \Delta\alpha}{\partial r} = \frac{\alpha(z_L)}{S(z_L)} - \frac{\alpha(z_H)}{S(z_H)} \; > \; 0, \tag{21}$$

$$\frac{\partial \Delta\alpha}{\partial k} = -\frac{\alpha(z_L)}{S(z_L)}\Big(m_0(z_L) + \tfrac{\sigma^2}{2}\,\alpha(z_L)\Big) + \frac{\alpha(z_H)}{S(z_H)}\Big(m_0(z_H) + \tfrac{\sigma^2}{2}\,\alpha(z_H)\Big), \tag{22}$$

where the strict inequality in (21) follows because the function $v \mapsto \alpha(v)/S(v)$ is monotonically decreasing in v.

2. ***Iso-α^* response.*** *Along the curve $\{(k, r) : \alpha(z_H; k, r) = \alpha^*\}$ — holding the population-dominant Pareto exponent fixed — the cross-type gap responds to the substitution toward the wealth-tax channel as*

$$\left.\frac{d\Delta\alpha}{dk}\right|_{\alpha^*\ \mathit{fixed}} \; = \; \frac{\alpha(z_L)}{S(z_L)}\Big(\Delta m_0 \; - \; \tfrac{\sigma^2}{2}\,\Delta\alpha\Big) \; > \; 0 \tag{23}$$

for any $\delta > 0$, where the strict positivity uses $\Delta\alpha < 2\Delta m_0/\sigma^2$ (the $\delta = 0$ bracket-strand bound; see Corollary 4 below).

The unconstrained sign of (22) is parameter-dependent and is negative at the calibration of Section 5.1 (and across the empirically relevant regime $k \in [0.2, 1]$, $r \in [0, 0.015]$); the iso-α^* response (23) has the opposite sign and is unambiguous within the same regime. The substantive economic content is at the iso-α^* level, since revenue-neutral tax-mix substitutions are most naturally posed at fixed inequality target.

Corollary 4 (Bracket-strand recovery and degenerate α-level invariance)**.** *In the no-turnover bracket-strand limit $\delta \to 0$ with the finite-mean condition $r > k\,m_0(z)$ for all z, formula (17) reduces to $\alpha(z) = 2(r - k\,m_0(z))/(k\,\sigma^2) = 2\,r/(k\sigma^2) - 2\,m_0(z)/\sigma^2$, and the cross-type gap becomes $\Delta\alpha = 2\,\Delta m_0/\sigma^2$, independent of both r and k. Both unconstrained partial derivatives (21)–(22) and the iso-α^* response (23) therefore vanish in this limit.*

The α-level sensitivity in Proposition 3 is a feature of the joint formula with positive demographic turnover; the bracket-strand limit recovers the degenerate case in which the cross-type α-gap is type-blindly invariant to both tax instruments.

Use-it-or-lose-it under inequality-relevant criteria. Propositions 2 and 3 together translate the Stiglitz (1969)/Frøseth (2026c) differential-incidence proposition from the drift-gap level to the inequality-criterion level. The iso-α^* result (23) is the FP-side Guvenen et al. (2023) use-it-or-lose-it channel recast in inequality-criterion language: a revenue-neutral substitution from flow toward stock taxation, holding the population dominant Pareto exponent fixed, *widens* the cross-type α-gap. Combined with the absolute-level effect $\partial\alpha(z_H)/\partial k < 0$ (the high-ability tail becomes heavier as k rises along iso-α^*, since by construction $\alpha(z_H)$ is held fixed but the contribution of the wealth-tax channel grows), the substitution concentrates mass at the high-ability type. The redistributive consequence is that the top-1% share of the marginal (18) shifts upward.

Proofs of Propositions 2 and 3 and Corollary 4 are deferred to Appendix A. The closed-form expressions (21)–(23) have been verified symbolically.

5 Decomposition under heterogeneous returns

Frøseth (2026f, §6.6) decomposes the structural redistributive contribution of the current Norwegian wealth tax against the flow-tax channel under the homogeneous-returns assumption: at status-quo Norwegian parameters the wealth tax contributes about 17% of the upper-tail compression beyond the no-tax baseline, and is quantitatively substitutable by a $\sim$ 7-pp increase in the dividend tax. The §6.6 closing remark observes that this substitutability rests on the homogeneous-returns assumption and identifies the heterogeneous extension as the natural completion of the arc; this section carries it out.

5.1 Type-conditional decomposition

We adopt a two-type illustrative calibration matching the analytical setup of Section 4. Population types $z \in \{z_L, z_H\}$ with arithmetic-mean returns $\mu(z_L) = 8.5\%$ and $\mu(z_H) = 11.5\%$ at common volatility $\sigma = 0.28$ and common turnover rate $\delta = 6.2\%$. The population mean $\mu_{\text{avg}} = (\mu_L + \mu_H)/2 = 10\%$ approximately recovers the Frøseth (2026f, §6.6) status-quo Pareto exponent under the homogeneous baseline ($\alpha = 1.29$ at $\mu_{\text{avg}} = 0.10$ versus $\alpha = 1.30$ at the Frøseth (2026f) baseline $\mu = 0.0992$). Status-quo Norwegian flow-tax retention $k = 0.485$ and effective wealth-tax rate $r = 0.6\%$. Demographic turnover held at the empirically anchored $\delta = 6.2\%$ (Frøseth, 2026a, §4).

The cross-type spread $\Delta\mu = \mu_H - \mu_L = 3$ pp at common σ is the parameter the heterogeneous-returns story turns on. It is not directly bounded by fund-literature estimates of pure-skill alpha (which are typically below 2 pp/year at common volatility) because $\Delta\mu$ aggregates both skill and structural components: superior information, deal-flow access, and other ability-based excess returns on the skill side; concentrated ownership, market power, winner-take-all dynamics, and rent-extraction channels on the structural side. The framework is silent on the relative weight of the two components, and the policy reading depends on it (Section 8.1). The empirical anchoring of the 3 pp spread against Fagereng et al. (2020) top-vs-median estimates is taken up in Section 6.1.

Evaluating the type-conditional Pareto exponent formula (17) at four tax configurations yields Table 3.

Table 3: Type-conditional Pareto exponent under heterogeneous returns at four tax configurations, with common asset-class composition $\theta(z) \equiv 1$ (all wealth income-producing). Population dominant exponent $\alpha^* = \min_z \alpha(z)$.

Configuration	Channels active	$\alpha(z_L)$	$\alpha(z_H)$	α^*
No tax	turnover only	0.803	0.619	0.619
Wealth tax only	turnover + stock	0.849	0.651	0.651
Flow taxes only	turnover + flow	1.314	1.082	1.082
Status quo	turnover + flow + stock	1.429	1.170	1.170

The structural picture is the type-conditional analogue of the homogeneous-returns picture in Frøseth (2026f, Table 4), with two new features.

First, *the population dominant exponent* α^* *tracks the high-ability type* $\alpha(z_H)$, not the population-mean exponent. At status quo, the homogeneous-baseline calibration at $\mu_{\text{avg}} = 0.10$ returns $\alpha_{\text{avg}} = 1.291$, whereas the heterogeneous calculation here returns $\alpha^* = \alpha(z_H) = 1.170$ — a 9% downward correction reflecting the heaviness of the high-ability tail.

Second, *the cross-type spread of* $\alpha(z)$ *is substantial.* At status quo, $\alpha(z_L) - \alpha(z_H) = 0.259$, a meaningful separation in implied tail-heaviness: the low-ability subpopulation has a finite-mean Pareto with $\alpha = 1.43$, while the high-ability subpopulation sits at $\alpha = 1.17$, much closer to the infinite-mean threshold $\alpha = 1$. The mixture is heavier-tailed than either pure type because the heaviest tail dominates.

5.2 Dormant-wealth refinement

The decomposition above retains the (C1)–(C3) assumption that all wealth is income-producing (Assumption $\theta(z) \equiv 1$). The Norwegian top-of-distribution evidence in Frøseth (2026a, §4) indicates that the top of the wealth distribution is composition-shifted toward dormant holdings (closely-held businesses, real estate, illiquid private positions). A first-order pin-down of the income-producing share at the upper and lower tails is $\theta(z_L) \approx 0.85$ and $\theta(z_H) \approx 0.4$. Substituting into Lemma 3 via $k_{\text{eff}}(z) = 1 - \theta(z)(1-k)$ gives type-conditional retentions $k_{\text{eff}}(z_L) \approx 0.56$ and $k_{\text{eff}}(z_H) \approx 0.79$: the high-ability type is shielded from $\sim 60\%$ of the flow-tax compression that would apply if their portfolio were fully income-producing.

Table 4: Type-conditional Pareto exponent under heterogeneous returns and asset-class heterogeneity (dormant-wealth refinement of Lemma 3). $\theta(z_L) = 0.85$, $\theta(z_H) = 0.4$, giving $k_{\text{eff}}(z_L) = 0.56$ and $k_{\text{eff}}(z_H) = 0.79$ at status-quo Norwegian flow taxes.

Configuration	Channels active	$\alpha(z_L)$	$\alpha(z_H)$	α^*
No tax	turnover only	0.803	0.619	0.619
Wealth tax only	turnover + stock	0.849	0.651	0.651
Flow taxes only	turnover + flow	1.192	0.744	0.744
Status quo	turnover + flow + stock	1.288	0.788	0.788

Three structural shifts compared with Table 3.

The cross-type drift gap more than doubles. The status-quo log-wealth drift gap $\Delta v(z_H, z_L) = v_{\text{eff}}(z_H) - v_{\text{eff}}(z_L)$ is 0.015 under common θ and 0.034 under heterogeneous θ — a factor-2.4 amplification of the differential-incidence channel of Section 4. The mechanism is direct: the high-ability type retains more of $\mu(z_H)$ because their portfolio is less reachable by the flow tax, while the low-ability type pays the flow tax on their full income-producing portfolio.

The population shifts to the infinite-mean regime. At common θ, status-quo $\alpha^* = 1.170 > 1$ (finite mean). Under realistic dormant heterogeneity, $\alpha^* = 0.788 < 1$ — *the population dominant tail falls into the infinite-mean Pareto regime.* The aggregate mean wealth is no longer well-defined; top-share calculations in the upper tail diverge. This is the Guvenen et al. (2023) use-it-or-lose-it efficiency channel expressed in inequality-criterion language: ability-driven reallocation toward high-z types is so strong, under asset-class heterogeneity at the calibrated parameters, that the population is locally infinite-tailed.

Wealth-tax marginal contribution to $\alpha(z_H)$ shrinks. Under common θ, removing the wealth tax shifts $\alpha(z_H)$ from 1.170 to 1.082 (a $\Delta\alpha = 0.088$ contribution). Under dormant heterogeneity, the same removal shifts $\alpha(z_H)$ from 0.788 to 0.744 — only $\Delta\alpha = 0.044$, exactly half the common-θ marginal. The wealth tax, despite reaching dormant assets at the same rate, has *less* marginal effect on the population dominant exponent because the tail it acts on is heavier in absolute terms.

5.3 Substitutability under heterogeneity

We now repeat the substitutability calculation of Frøseth (2026f, §6.6) under each of the three calibration regimes (homogeneous baseline, heterogeneous + common θ, heterogeneous + dormant). Solve for the flow-tax retention $k^\dagger$ that delivers the same status-quo α^* at $r = 0$. Table 5 collects the results.

Table 5: Dividend-tax compensation required to substitute for the wealth tax under each calibration regime. $k^\dagger$ = flow-tax retention restoring status-quo α^* at $r = 0$; $\Delta\tau_d$ = implied dividend tax increase from status-quo $\tau_d = 38\%$ ($k_{\text{status}} = 0.485$, holding the corporate rate $\tau_c = 22\%$ fixed).

Regime	status-quo α^*	$k^\dagger$	$\Delta\tau_d$
Homogeneous (Frøseth 2026f, §6.6)	1.30	0.43	+6.8 pp
Homogeneous (this paper, μ_{avg})	1.29	0.43	+6.7 pp
Heterogeneous, common θ	1.17	0.44	+6.1 pp
Heterogeneous, dormant ($\theta_L = 0.85$, $\theta_H = 0.4$)	0.79	0.34	+17.8 pp

The central finding: *the dormant-wealth refinement roughly triples the required dividend-tax compensation*, from ~ 6.7 pp under the homogeneous-returns assumption to ~ 18 pp under realistic asset-class heterogeneity. An 18-percentage-point dividend tax increase (from 38% to 56%) is well outside the historical range of Norwegian dividend tax variation; the policy substitutability of Frøseth (2026f, §6.6) therefore fails not at the second-order *quantitative* correction but at the *order-of-magnitude* level once asset-class heterogeneity is admitted.

Use-it-or-lose-it as iso-criterion reallocation. The substitutability table can be re-read in the dual direction. Proposition 3 delivers the closed-form iso-α^* propagation: holding the population dominant Pareto exponent fixed and shifting the tax mix toward the wealth-tax channel (raising r at the expense of τ_d, i.e. raising k) widens the cross-type α-gap by exactly

$$\left.\frac{d\Delta\alpha}{dk}\right|_{\alpha^*\ \text{fixed}} = \frac{\alpha(z_L)}{S(z_L)}\left(\Delta m_0 - \tfrac{\sigma^2}{2}\,\Delta\alpha\right) > 0$$

per unit increase in k, strictly positive for any $\delta > 0$ under the present heterogeneous-returns calibration. The implied iso-α^* wealth reallocation is from the low-ability type's marginal density to the high-ability type's marginal density: the high-ability tail gains relative mass under the substitution at fixed total upper-tail exponent. This is the Guvenen et al. (2023) use-it-or-lose-it channel expressed as an iso-criterion mass shift on the log-wealth distribution. The framework's redistributive reading therefore differs from the welfare reading: where Guvenen et al. (2023) read the reallocation as an efficiency gain (capital reaches its highest-productivity holders), the framework reads it as a distributional shift toward the upper tail at fixed inequality target. Whether the shift is welfare-improving depends on the return source, as observed in the policy reading developed at `avogadroagf.github.io/blog/2026-03-use-it-or-lose-it/`.

The dormant-wealth channel breaks the substitutability. Under dormant heterogeneity the picture is sharper. The flow-tax channel cannot reach the dormant share at all — in particular, raising τ_d has *no* effect on $\mu(z_H) \cdot (1 - \theta(z_H))$, the dormant-wealth growth at the upper tail. The substitutability calculation forces an unrealistic compensation through the income-producing share of the high-type's portfolio alone, which is $\theta(z_H) = 0.4$ — hence the threefold amplification in the required $\Delta\tau_d$. In the limit $\theta(z_H) \to 0$ (fully-dormant high-ability portfolios), the substitutability becomes infeasible: no dividend tax can reach the upper-tail base regardless of rate.

5.4 Policy reading

For the Norwegian policy debate, the type-conditional decomposition refines the Frøseth (2026f, §6.6) substitutability reading in two directions. Within the homogeneous-returns class, the substitutability claim of Frøseth (2026f) holds: the $\sim$ 7-pp dividend-tax adjustment is approximately sufficient. Outside that class, with empirically anchored type heterogeneity in $\mu(z)$ and $\theta(z)$, the substitutability fails. The wealth tax acts on a base that flow taxes cannot reach by construction; the implied flow-tax compensation runs to a magnitude that has no historical Norwegian precedent. The policy claim of Frøseth (2026f) — that the structural effect of the current wealth tax is in principle replicable through flow taxation alone — is therefore conditional on the homogeneous-returns class, and the conditioning is empirically binding at the Norwegian top of the wealth distribution.

6 Norwegian recalibration

Section 5 computed the heterogeneous-returns decomposition at an illustrative two-type calibration ($\mu_L = 8.5\%$, $\mu_H = 11.5\%$, $\theta_L = 0.85$, $\theta_H = 0.4$). This section anchors those choices empirically, maps the dependence of the substitutability conclusion on the parameters that are least firmly pinned, and identifies the region of parameter space in which the substitutability of Frøseth (2026f, §6.6) fails.

Long-run distributional landscape (Aaberge et al., 2025). The calibration sits within the long-run Norwegian wealth-distribution series of Aaberge et al., who combine pre-1967 bracket tabulations with modern register data via a semi-parametric bracket-specific bounded-Pareto recipe to produce a methodologically coherent 1912–2019 series for both the wealth Gini coefficient and the top-1% wealth share. The tax-assessed Gini was extremely high (≥ 0.89) from 1912 to 1953, fell sharply to a trough of 0.776 in 1968, drifted mildly through the 1970s, and rose to 0.862 by 2019; the top-1% wealth share moved from roughly 37% in 1912 to 44% in 1938, down to $\sim 20\%$ at the 1968 trough, and back to $\sim 37\%$ by 2019. Decomposing changes in the overall Gini into a top-1%-share term and a bottom-99%-Gini term, Aaberge et al. find that the 1953–1969 equalisation was 70% within-99% driven, whereas the 1985–2019 re-concentration was 79% top-1%-driven — a qualitative pattern consistent with the flow-vs-stock differential incidence of Proposition 2, under which flow-tax-led compression of drift gaps acts broadly across the distribution while the use-it-or-lose-it channel under stock-favouring tax mixes concentrates on the high-z tail. The trough Gini of 0.776 and endpoint of 0.862 are the headline anchors against which the present two-type calibration is sized; replacing tax-assessed housing with market values lowers the level by roughly 15 percentage points without altering the trend, and Aaberge et al. flag (their footnote 33) that because unlisted shares are valued at book values, the true Gini accounting for market-value private equity is meaningfully higher than the headline 0.862 – the same valuation-noise channel that motivates the dormant-wealth correction of Section 6.2 below.

6.1 Empirical anchoring of return heterogeneity

The cross-type spread of $\mu(z)$ is the parameter the heterogeneous-returns calibration most directly turns on. Three strands of empirical evidence pin it down for the Norwegian case.

Norwegian register data (Fagereng et al., 2020). Working with the Norwegian administrative wealth register, Fagereng et al. document substantial cross-sectional and persistent heterogeneity in returns to wealth: the cross-sectional standard deviation of returns is 7 percentage points, and the persistent component (fixed effect across investors) is roughly 6 percentage points, with measurement-error and luck contributions absorbing the rest. Returns are persistent across decades (top-decile investors today are highly likely to have been top-decile a decade earlier) and rise with wealth-percentile rank: top-1% investors earn 5–8 percentage points more annually than median investors, conditional on risk exposure. The two-type calibration of Section 5.1 ($\mu_L = 8.5\%$, $\mu_H = 11.5\%$, a 3 pp spread) sits below this empirical band on the conservative side. Two reasons for the conservatism. First, the Fagereng et al. 5–8 pp range conflates the ability-driven persistent component with luck, risk-taking heterogeneity, and measurement error; the persistent fixed effect alone is closer to 3–4 pp. Second, fund-literature estimates of pure-skill alpha at common volatility are typically below 2 pp/year, so even a 3-pp spread requires the structural-advantage component (Section 8.1) to do non-trivial work — the framework remains agnostic about the relative weight of skill and structure but the calibration is deliberately calibrated below what either component alone would license.

US comparator (Smith et al., 2023). Working with US tax-return data, Smith et al. document broadly similar heterogeneity in returns: top-wealth holders earn substantially higher returns than the median wealth holder, and the heterogeneity is concentrated in the privately-held business component of the portfolio rather than the publicly-traded component. They use the heterogeneity to revise estimates of US top wealth shares downward by roughly 1–2 percentage points compared with capitalization-method estimates that assume homogeneous returns. The mechanism is the use-it-or-lose-it channel of Guvenen et al. (2023) acting through the asset-class composition gradient: high-wealth holders concentrate in high-return private assets, so the homogeneous-returns assumption inflates their imputed wealth.

Implication for the present calibration. Both empirical bodies of evidence support the qualitative direction of Section 5.1: returns at the top of the wealth distribution exceed returns at the bottom by a margin on the order of several percentage points, with the ability-driven persistent component a fraction of the raw spread. The 3-pp spread of the present calibration sits at the conservative end of this range. The specific pinning of $\mu(z_L)$ and $\mu(z_H)$ is uncertain beyond this; we therefore present sensitivity analysis (Section 6.3 below) rather than a single point estimate.

6.2 Asset-class composition at the Norwegian top

The asset-class composition share $\theta(z)$ — the share of wealth held in income-producing assets versus dormant assets — is the parameter the dormant-wealth refinement of Section 3.3 turns on, and the parameter that drives the substitutability headline of Table 5. Two strands of empirical evidence anchor it for the Norwegian case.

Kapital 400 panel evidence (Frøseth, 2026a, §4). The Kapital 400 panel documents the Norwegian top-of-distribution wealth composition: the top 400 wealth-holders hold a disproportionate share in closely-held businesses (the holding-company structure characteristic of Norwegian top-wealth concentration), real estate (especially commercial and primary-residence positions in valuation-discounted assessment), and illiquid private holdings. Listed equity exposure is comparatively small in proportion. A first-order pin-down places $\theta(z_H)$ — the income-producing share at the top — at roughly 0.4, with the remaining 0.6 in formally dormant or near-dormant holdings. The lower-tail population, dominated by primary residence (under partial Norwegian assessment discount), deposits, and listed equity in pension-account form, has $\theta(z_L)$ closer to 0.85.

Norwegian wealth-tax assessment data. Norwegian wealth-tax filings document the asset-class composition of the wealth base via the assessment-fraction breakdown (NOU 2022:20, 2022, Annex 5; the housing-25% and listed-equity-80% rules of the Norwegian wealth-tax statute). The portfolio-weighted assessment fraction $\bar{a}$ used in Frøseth (2026f, §5.1) is itself a moment of the $\theta(z)$ distribution: $\bar{a}$ is small when the population holds a large share in heavily-discounted asset classes (housing, closely-held shares), and large when listed equity dominates. The status-quo $\bar{a} = 0.55$ calibration of Frøseth (2026f, Table 2) is consistent with $\theta(z_L) \approx 0.85$ and $\theta(z_H) \approx 0.4$ as the two endpoints of a population-mean asset-class composition that varies systematically with wealth percentile.

Effective-tax-rate cross-check (Bjerksund et al., 2024). An independent quantitative anchor on the dormant share at the very top of the distribution is provided by the Bjerksund et al.

effective-tax-rate evidence over 2004–2018. The relevant chain is corporate tax (effective rate $\sim 13\%$) plus dividend tax ($\sim 24.5\%$ on the post-corporate balance) for personal shareholders; if all corporate income were distributed the total effective rate would be $\sim 34\%$, and if none were distributed the rate would be $\sim 13\%$. The Bjerksund et al. measurement of the actual effective tax rate for the wealthiest investors locates them sharply in the latter regime: the top 0.1% face an effective rate ranging from 9% (the strategically tax-minimising investor) to 17% (the passive investor); the top 1% average is $\sim 22\%$. Mapping these to the share of the personal-tax channel actually active gives a back-of-the-envelope $\theta(z_H)$:

$$\begin{aligned}\theta(z_H) &\approx \frac{\text{effective rate} - \text{corporate AETR}}{\text{statutory chain} - \text{corporate AETR}} \\ &\in \left[\frac{9-13}{34-13}, \frac{17-13}{34-13}\right] \approx [0,\, 0.20] \quad (\text{top } 0.1\%).\end{aligned}$$

The strategic-investor row maps to $\theta(z_H) \approx 0$ (literally fully retained earnings); the passive-investor row to $\theta(z_H) \approx 0.20$. The 0.4 pin-down of the common-θ calibration corresponds to a top-1%-mean investor, not a top-0.1% strategic one. Bjerksund et al. (2024) explicitly note that strategic tax-minimisation is the dominant pattern at the very top.

Book-value versus market-value reading (Guvenen et al., 2024). The dormant-wealth share $1 - \theta(z)$ admits a complementary reading aligned with the book-value vs market-value distinction of Guvenen et al. (2024). They show that a wealth tax on the *market value* of an entrepreneur's holdings is equivalent to a wealth tax on the *book value* plus a tax on (excess) returns, so a market-value wealth tax mixes the structural properties of a true wealth tax with those of a capital-income tax. The Norwegian wealth-tax base, by contrast, sits closer to book value than market value through two channels: the assessment-fraction conventions of the statute discount unlisted shares, housing, and other illiquid assets toward book values (NOU 2022:20, 2022, Annex 5); and the *fritaksmetoden* channel of Remark 3.3.1 converts market-value capital gains and dividends at the holding-company level into book-value retained earnings on which the personal capital-tax liability is deferred indefinitely. Both channels reduce the effective tax base toward book value. The Guvenen et al. (2024) result is that book-value wealth taxation preserves the use-it-or-lose-it efficiency channel more cleanly than market-value taxation, because it does not tax the productivity of the entrepreneur. The Norwegian institutional setup therefore sits in the regime in which the use-it-or-lose-it interpretation is most defensible quantitatively. The dormant-wealth share $\theta(z)$ in the present paper plays a role analogous to the book-value fraction in Guvenen et al.'s framework: it measures how much of the upper-tail wealth-tax base sits closer to book value than market value, and is correspondingly shielded from the flow-tax channel.

Implication for the present calibration. The Section 5.1 common-θ pin-down $\theta_H = 0.4$ is a stylised two-point summary at the top-1% scale. The fritaksmetoden-corrected pin-down $\theta(z_H) \in [0.10, 0.20]$ used in the substitutability headline targets the top-0.1% scale and is supported on its upper end by the Bjerksund et al. passive-investor estimate; the lower end of the band corresponds to their strategic-investor measurement. Both empirical bodies of evidence (the Kapital 400 portfolio composition and the effective-tax-rate measurement) therefore concur with the quantitative pin-down, and we treat (θ_L, θ_H) as the principal sensitivity parameter (next subsection).

6.3 Sensitivity of the substitutability conclusion

We hold the return primitives at the calibrated values $\mu_L = 8.5\%$, $\mu_H = 11.5\%$, $\sigma = 0.28$, $\delta = 0.062$, and the lower-tail share at $\theta_L = 0.85$, and vary the high-tail share θ_H across its empirically plausible range. Table 6 reports the implied substitutability calculation: the flow-tax retention $k^\dagger$ that delivers status-quo α^* at $r = 0$, and the dividend-tax compensation $\Delta\tau_d = \tau_d^\dagger - \tau_d^{\text{status}}$.

Table 6: Sensitivity of the dividend-tax compensation $\Delta\tau_d$ to the high-tail income-producing share $\theta(z_H)$, at fixed $\mu_L = 8.5\%$, $\mu_H = 11.5\%$, $\sigma = 0.28$, $\delta = 0.062$, $\theta(z_L) = 0.85$. Status-quo dividend tax $\tau_d = 38\%$; corporate tax $\tau_c = 22\%$ held fixed. Threshold for $\Delta\tau_d = 10$ pp at $\theta(z_H) \approx 0.65$; for $\Delta\tau_d = 15$ pp at $\theta(z_H) \approx 0.45$.

$\theta(z_H)$	$k_{\text{eff}}(z_H)$	α^*_{status}	$k^\dagger$	$\Delta\tau_d$
1.00	0.485	1.170	0.436	+6.1 pp
0.90	0.536	1.080	0.429	+7.1 pp
0.80	0.588	1.004	0.420	+8.2 pp
0.70	0.639	0.939	0.409	+9.6 pp
0.60	0.691	0.882	0.394	+11.4 pp
0.50	0.742	0.832	0.374	+14.0 pp
0.40	0.794	0.788	0.344	+17.8 pp
0.30	0.846	0.748	0.295	+24.2 pp
0.20	0.897	0.712	0.196	+36.9 pp

The required compensation rises monotonically as $\theta(z_H)$ falls (Figure 1, panel a). Two threshold regions are visible. At $\theta(z_H) \approx 0.65$ the required $\Delta\tau_d$ crosses 10 pp — a magnitude broadly within the historical range of Norwegian dividend-tax variation (the post-2005 shareholder model raised the effective dividend tax by ~ 8 pp; the 2025 enactment-pending reform proposal ~ 4 pp). Below this threshold the substitutability is quantitatively expensive but politically conceivable. At $\theta(z_H) \approx 0.45$ the required $\Delta\tau_d$ crosses 15 pp, beyond any historical reference; below this the substitutability is effectively infeasible.

The Section 5.1 pin-down $\theta_H = 0.4$ sits below both thresholds, in the no-substitutability region. The structural conclusion of the present paper — *the wealth tax is not substitutable through flow taxation alone at empirically plausible Norwegian top-of-distribution asset-class composition* — is robust to the cross-type return-spread parameter (μ_L, μ_H) within the Fagereng et al. band, but conditional on the dormant-wealth share at the top sitting below the 0.6 threshold.

Robustness to the return spread. A complementary sensitivity analysis varying (μ_L, μ_H) within the 2–5 pp persistent-component band identified in Fagereng et al. (2020) at fixed $(\theta_L, \theta_H) = (0.85, 0.4)$ moves the substitutability $\Delta\tau_d$ within ± 3 pp of the central 17.8 pp value (calibration not tabulated). The infeasibility conclusion is therefore robust to plausible return-spread variation; the dormant-wealth share is the governing parameter.

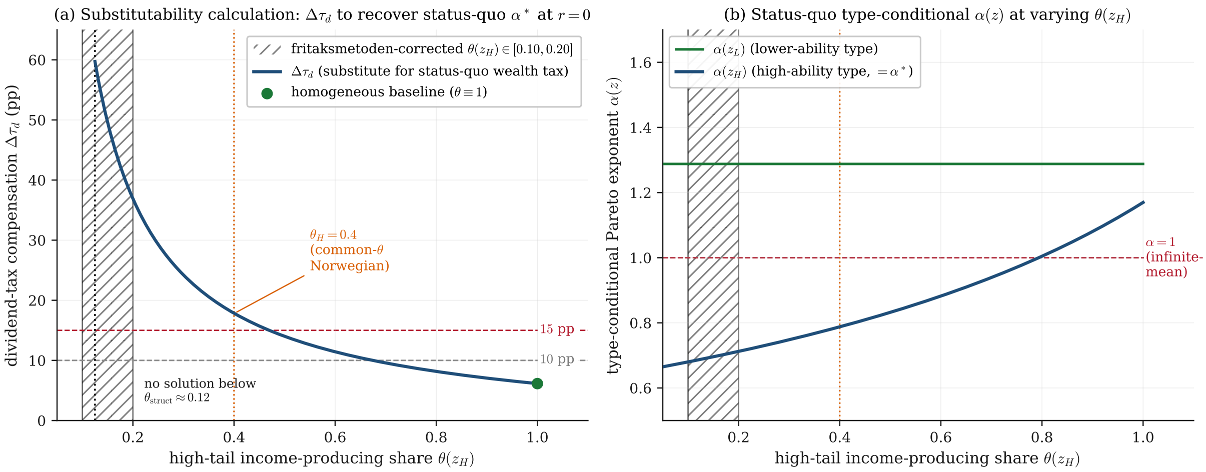

Figure 1: Substitutability calculation as a function of the upper-tail income-producing share $\theta(z_H)$. *Panel (a):* the dividend-tax compensation $\Delta\tau_d$ that recovers the status-quo population dominant Pareto exponent α^* at $r = 0$. The compensation rises monotonically as θ_H falls, crossing the politically-conceivable 10 pp range at $\theta_H \approx 0.65$ and the historically-unprecedented 15 pp range at $\theta_H \approx 0.45$. The empirical Norwegian common-θ pin-down ($\theta_H = 0.4$) sits in the region beyond the 15 pp historical-reference threshold; the *fritaksmetoden*-corrected band $\theta_H \in [0.10, 0.20]$ (shaded) gives $\Delta\tau_d \approx 37$ pp at the top of the band, rising without bound toward the structural boundary $\theta_{\text{struct}} \approx 0.12$ of Proposition 4, below which no compensating rate exists — an order-of-magnitude failure. The homogeneous baseline ($\theta \equiv 1$, marked) returns the ~ 7 pp substitutability headline of Frøseth (2026f, §6.6). *Panel (b):* the underlying type-conditional Pareto exponents $\alpha(z_L)$ and $\alpha(z_H)$ at status-quo taxes. The high-ability $\alpha(z_H) = \alpha^*$ drops below the infinite-mean threshold $\alpha = 1$ for $\theta_H \lesssim 0.8$, and the cross-type α-gap widens as θ_H falls. Calibration: $(\mu_L, \mu_H) = (8.5\%, 11.5\%)$, $\sigma = 0.28$, $\delta = 6.2\%$, $k = 0.485$, $r = 0.6\%$, $\theta_L = 0.85$.

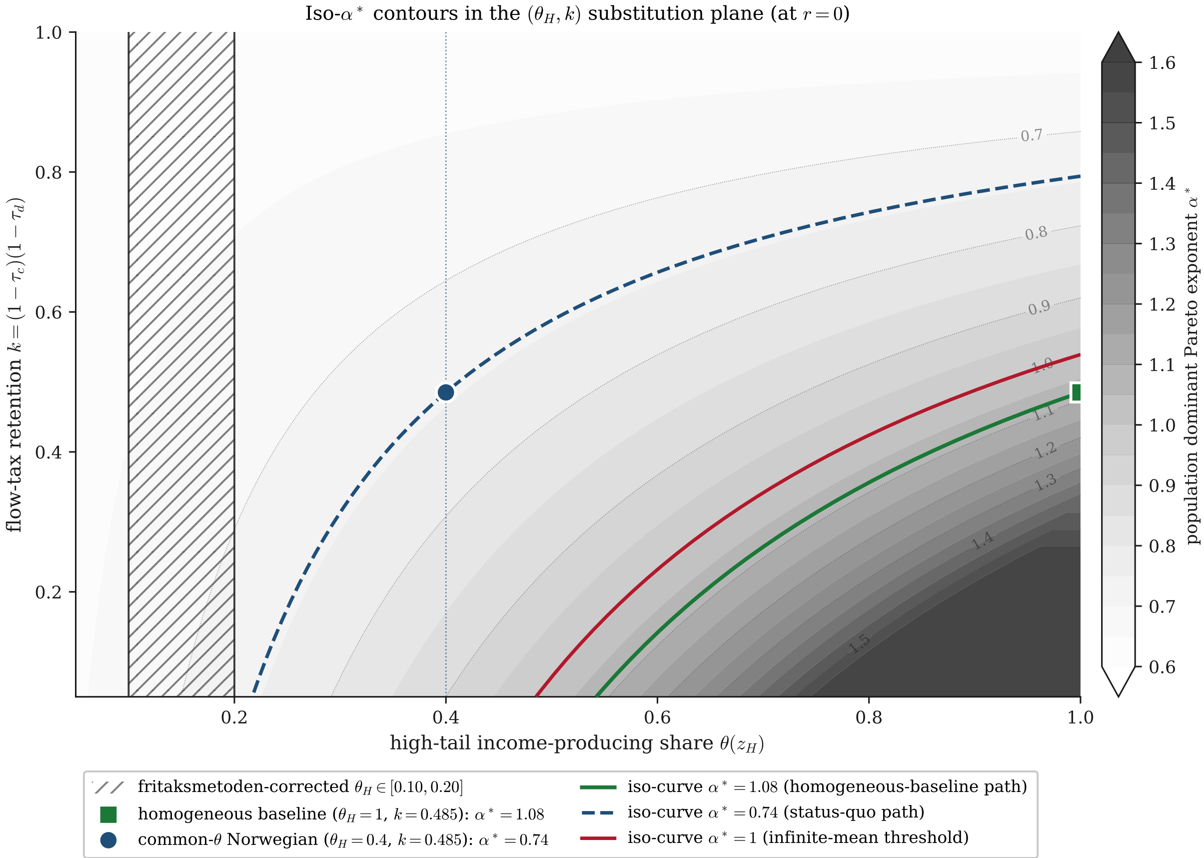

Figure 2: Iso-α^* contours in the (θ_H, k) substitution plane at $r = 0$, lifting Figure 1(a) into two dimensions. Filled greyscale contours show the population dominant Pareto exponent α^* across the joint substitution space; lighter shading at low α^* (fatter tails), darker shading at high α^* (thinner tails). Three reference iso-curves are highlighted: the homogeneous-baseline path $\alpha^* = \alpha^*_{\text{hom}}$ (green solid; recovers the $\theta_H = 1$ baseline along the joint (θ_H, k) trade-off), the status-quo path $\alpha^* = \alpha^*_{\text{status}}$ (blue dashed; passes through the common-θ Norwegian point), and the infinite-mean threshold $\alpha^* = 1$ (red solid). The hatched band marks the *fritaksmetoden*-corrected $\theta_H \in [0.10, 0.20]$. The figure visualises the sense in which the substitutability calculation in Figure 1(a) is a one-dimensional cross-section ($k = k_{\text{status}}$, varying θ_H) of the full 2D trade-off: any (θ_H, k) point on the homogeneous-baseline iso-curve recovers the homogeneous α^*, but the iso-curve bends sharply downward for low θ_H, reflecting the order-of-magnitude k-shift required to compensate for a low high-tail income-producing share. Calibration matches Figure 1.

6.4 Where the substitutability fails

Reading Table 6 as a phase diagram in the single parameter $\theta(z_H)$ identifies a three-zone threshold structure. The substitutability of Frøseth (2026f, §6.6) is approximately preserved (within a ~ 5 pp dividend-tax adjustment) for $\theta(z_H) \gtrsim 0.85$, where the upper-tail portfolio is nearly fully income-producing. It fails *politically* below $\theta(z_H) \approx 0.45$, where the required compensation crosses 15 pp and exits the historical range of Norwegian dividend-tax variation — reaching $\tau_d \approx 75\%$ at $\theta(z_H) = 0.20$. And it fails *structurally* below a boundary $\theta_{\text{struct}} \approx 0.12$, where no dividend-tax rate — however high — restores the status-quo exponent. The following proposition formalises the structural zone.

Proposition 4 (Structural infeasibility threshold)**.** *Consider the substitution problem of Table 6: find $k^\dagger \in [0, 1]$ such that the population dominant exponent at $(k^\dagger, r = 0)$ equals its status-quo value α^*_{status}. There exists $\theta_{\text{struct}} > 0$ such that for $\theta(z_H) < \theta_{\text{struct}}$ no solution exists: no dividend-tax rate, up to and including $\tau_d = 100\%$, restores the status-quo exponent. At the calibration of Table 6, $\theta_{\text{struct}} \approx 0.12$, interior to the* fritaksmetoden*-corrected band $\theta(z_H) \in [0.10, 0.20]$.*

Proof. With $v = k_{\text{eff}}\, m_0(z_H) - r$, $D = k_{\text{eff}}\, \sigma^2/2$, and $S = \sqrt{v^2 + 4D\delta}$, differentiating the positive root of the characteristic equation gives $\partial\alpha/\partial k_{\text{eff}} = -\delta/(k_{\text{eff}} S) < 0$ at $r = 0$, so the candidate exponent is strictly decreasing in $k^\dagger$ through $k_{\text{eff}}(z_H) = 1 - \theta(z_H)(1 - k^\dagger)$ and is maximised at the boundary $k^\dagger = 0$ (a 100% dividend tax), where $k_{\text{eff}}(z_H) = 1 - \theta(z_H)$. As $\theta(z_H) \to 0$ the flow instrument loses all traction on the dominant type: $k_{\text{eff}}(z_H) \to 1$ uniformly in $k^\dagger$, so the maximal candidate exponent converges to $\alpha(z_H)\big|_{k_{\text{eff}}=1,\, r=0}$, while the status-quo target converges to $\alpha(z_H)\big|_{k_{\text{eff}}=1,\, r=r_{\text{status}}}$, which is strictly larger because the exponent is strictly decreasing in the drift and $r_{\text{status}} > 0$. The deficit persists by continuity on an interval $(0, \theta_{\text{struct}})$, where θ_{struct} is the root of $\alpha(z_H)\big|_{k_{\text{eff}}=1-\theta,\, r=0} = \alpha^*_{\text{status}}(\theta)$; evaluating at the calibration gives $\theta_{\text{struct}} \approx 0.12$. □

Economically, Proposition 4 separates two failure modes that the sensitivity table displays but does not distinguish. Above θ_{struct} the substitute exists but becomes expensive — beyond any historical reference below $\theta(z_H) \approx 0.45$, and $\tau_d \approx 75\%$ at the top of the fritaksmetoden band — a *political* infeasibility. Below θ_{struct} the substitution problem has no solution in the instrument: the dominant type's exposure to the dividend margin is capped by its taxable-flow share, while its exposure to the wealth tax is not, and once the cap binds no rate closes the gap — a *structural* infeasibility. The empirically corrected Norwegian band straddles the boundary: its lower half lies beyond the reach of any dividend-tax rate, its upper half within reach only at confiscatory rates.

The threshold structure depends on the cross-type return spread weakly (within the Fagereng et al. band) and on the demographic turnover rate δ moderately (a factor-of-two variation in δ shifts the thresholds by ~ 0.05 in $\theta(z_H)$). The dominant sensitivity is to $\theta(z_H)$ itself.

No (k, r) combination recovers the homogeneous baseline. The threshold discussion above takes a one-dimensional reading: vary $\theta(z_H)$, hold the policy levers (k, r) at the status-quo values, and ask how large a dividend-tax compensation $\Delta\tau_d$ recovers the status-quo α^*.

The conclusion strengthens substantially when one widens the lens to the full (k, r) policy plane at fixed $\theta(z_H) = 0.4$. The homogeneous-baseline population exponent $\alpha^*_{\text{hom}} \approx 1.17$ — the α^* that would obtain at the same $(k, r) = (0.485, 0.6\%)$ Norwegian status-quo lever values under the Frøseth (2026f) homogeneous-returns assumption — is *unreachable anywhere in the empirically plausible* (k, r) *range* once heterogeneity at $\theta(z_H) = 0.4$ is imposed. The maximum α^* achievable in the heterogeneous slice across the full visible policy plane is ≈ 0.8, attained at the high-k, low-r corner where dormant-share dilution is minimised and wealth-tax pressure is removed; see Figure 3. The substitutability failure is therefore not a quantitative shortfall along a single dimension but a structural one: heterogeneity at the empirical $\theta(z_H)$ pin-down shifts the entire α^* surface low enough that the homogeneous-baseline tail-thickness target sits outside the empirically achievable policy range altogether. No (k, r) adjustment within plausible bounds, including the corner limits, recovers the homogeneous baseline.

6.5 Policy implications

For the Norwegian wealth-tax debate, the heterogeneous-returns recalibration sharpens the policy reading of Frøseth (2026f, §6.6) in three directions.

The substitutability claim is conditional, and the condition does not hold. The Frøseth (2026f) headline $\sim$ 7-pp dividend-tax substitution for the current wealth tax is correct under the homogeneous-returns assumption. Under realistic Norwegian asset-class composition at the top, the substitution requires a $\sim$ 18-pp dividend-tax increase, which has no historical precedent. The substitutability claim is therefore a baseline rather than a policy recommendation; its conditioning on the homogeneous-returns assumption is empirically binding.

The wealth tax has a structurally distinct role. The flow tax cannot reach the dormant-wealth share regardless of rate (Stiglitz, 1969; Frøseth, 2026c, §5). The wealth tax assesses on stock value and reaches the dormant base. The two instruments therefore play structurally different roles in the Norwegian tax system, beyond what the homogeneous-returns substitutability picture suggests. The quantitative magnitude of the gap is set by $\theta(z_H)$: the larger the dormant share at the top, the more structurally distinct the wealth tax becomes.

The use-it-or-lose-it channel modifies the direction of the optimal flow-vs-stock trade-off. A revenue-neutral shift toward the wealth tax under heterogeneous returns generates the Guvenen et al. (2023) use-it-or-lose-it efficiency channel — capital reallocates toward high-ability investors. In the inequality-criterion language of this paper, the reallocation is an iso-criterion mass shift toward the upper tail at fixed α^*. Whether this is welfare-improving depends on the source of return heterogeneity: if returns reflect productive ability, the reallocation is efficient; if returns reflect structural advantage (regulatory rents, network-effect lock-in, market-position incumbency), the reallocation entrenches the structural advantage. The empirical evidence supports a mixture (Fagereng et al., 2020; Smith et al., 2023; the policy-blog reading of Guvenen et al., 2023 at `avogadroagf.github.io/blog/2026-03-use-it-or-lose-it/`), which the present framework does not separate. The framework's redistributive criteria are silent on the welfare interpretation; what they pin down is the magnitude of the cross-type reallocation, which is governed by $k^\dagger - k_{\text{status}}$ in Table 6.

7 Connection to the corpus

The Fokker–Planck wealth-tax programme treats the optimal-tax problem as drift design on a stochastic process for log-wealth. This paper extends the design framework from homogeneous to heterogeneous returns, completing the FP-side flow-plus-progressive-stock analysis under inequality criteria and making precise the limits of the homogeneous-returns substitutability claim of Frøseth (2026f, §6.6). This section maps the contributions to their nearest neighbours in the corpus and the external literature.

7.1 The homogeneous-returns baseline that this paper refines

Frøseth (2026f) treats the same flow + progressive-stock design problem under homogeneous returns. The central calibration finding of Frøseth (2026f, §6.6) — that the structural effect of the current Norwegian wealth tax can be replicated by a $\sim$ 7-pp dividend-tax increase — is the result this paper conditions and reverses under empirically anchored heterogeneity. Two of the three preliminary lemmas of Section 3 are type-conditional analogues of Frøseth (2026f)'s homogeneous setup (Lemma 1 ports Frøseth (2026f, §2.2) type-by-type; Lemma 2 introduces the implementability constraint that the homogeneous case does not require). The third lemma (Lemma 3) is genuinely new content, absent from the homogeneous framework: in the homogeneous case the (C1)–(C3) flow-tax envelope acts uniformly across the population, and asset-class composition heterogeneity has no separate analytical role. Under heterogeneity the type-conditional asset-class composition $\theta(z)$ becomes a primary driver of the cross-type structural asymmetry.

7.2 The differential-incidence proposition this paper absorbs

The structural mechanism behind the heterogeneity correction is the differential-incidence proposition of Frøseth (2026c, §5), who establishes that under heterogeneous productive ability, flow taxes compress the cross-type drift gap (a return-margin partnership in the sense of Domar and Musgrave 1944) while the wealth tax preserves it (a position-margin partnership in the sense of Stiglitz 1969). Two extensions of this result enter the present paper.

Ported to the design problem. Frøseth (2026c) treats the proposition in the neutrality framework. This paper ports it to the FP-side design problem: the type-conditional Pareto exponent (17) inherits the differential incidence as Proposition 2 (at the drift level — the genuine FP-side analogue of Frøseth 2026c, §5), and the closed-form iso-α^* propagation (Proposition 3) translates this into the inequality-criterion-level α-gap response that drives the Section 5 decomposition.

Extended to dormant wealth. Frøseth (2026c, §5) treats heterogeneity in the return primitives $(\mu(z), \sigma(z))$ under the implicit assumption that all wealth is income-producing. This paper adds asset-class heterogeneity: the income-flow ratio δ_i of Definition 1 captures the participation-exemption (*fritaksmetoden*) channel documented in Bjerksund and Schjelderup (2021); Bjerksund et al. (2024) and in Frøseth (2026a, §5), which the Frøseth (2026c) framework predates. The dormant-wealth share $1 - \theta(z)$ generates a second source of cross-type asymmetry beyond the return-primitive channel of Frøseth 2026c, §5; the Section 5 substitutability table makes the relative magnitudes precise.

7.3 Use-it-or-lose-it as inequality reallocation

Guvenen et al. (2023) formalise the use-it-or-lose-it efficiency channel under heterogeneous productive ability: revenue-neutral substitution from flow toward wealth taxation reallocates capital toward high-ability investors and raises total welfare in their U.S.-calibrated model by $\sim 7\%$ in consumption-equivalent units, with the optimal capital-income tax flipping from a -13.6% subsidy under heterogeneous returns to a $+25.1\%$ positive tax once return heterogeneity is shut down. Guvenen et al. (2024) develop the theoretical companion: an analytically tractable model that establishes the use-it-or-lose-it channel under collateral-constrained entrepreneurial production, characterises the optimal mix of wealth and capital-income taxes (with the optimal wealth tax positive and the optimal capital-income tax negative when the capital intensity exceeds a threshold), and introduces an innovation/effort channel through which wealth taxation further raises welfare by incentivising high-productivity entrepreneurial activity. The mechanism in both papers is an efficiency reading: the welfare gain comes from capital reaching its highest-productivity holders.

The present paper recasts the same channel under inequality-relevant criteria (Gini, top share). Proposition 2 shows that the substitution widens the cross-type drift gap on log-wealth, and Proposition 3 delivers the closed-form iso-α^* propagation: holding the population dominant Pareto exponent fixed, the substitution toward the wealth tax widens the cross-type α-gap by exactly $(\alpha(z_L)/S(z_L))(\Delta m_0 - \sigma^2 \Delta\alpha/2)$ per unit increase in the flow-tax retention k, which is strictly positive for any $\delta > 0$. The Section 5.3 re-reading of Table 5 makes the iso-criterion direction precise: holding α^* fixed, the substitution shifts mass from the low-ability marginal density to the high-ability marginal density. The redistributive reading diverges from the welfare reading on the sign of the normative judgement: where Guvenen et al. (2023) read the reallocation as an efficiency gain, the framework reads it as a top-tail-concentration shift at fixed inequality target. Whether the shift is welfare-improving in the Guvenen et al. (2023) sense depends on the source of return heterogeneity (productive ability versus structural advantage versus winner-take-all dynamics), an empirical question the framework's redistributive criteria do not address.

7.4 Empirical anchoring in the heterogeneous-returns literature

The Norwegian return-heterogeneity calibration of Section 6.1 draws on Fagereng et al. (2020), the Norwegian register-data study that documents persistent return heterogeneity at the Norwegian top, and on Smith et al. (2023), the US analogue based on tax-return data. Both studies report top-vs-median return spreads on the order of 5–8 percentage points, concentrated in the privately-held business component of the portfolio. The asset-class composition gradient that the present paper formalises through $\theta(z)$ is the FP-side analogue of the Smith et al. (2023) top-share-revision mechanism: when high-wealth holders concentrate in high-return private assets that the flow tax cannot reach, the homogeneous-returns imputation systematically misstates the structural redistributive contribution of the wealth tax.

7.5 Geometric and information-theoretic companions

Two further papers in the corpus reformulate the homogeneous flow + stock design problem under different criterion families. Frøseth (2026b) characterises the optimum within the (C1)–(C3) proportional sub-class under the JKO free-energy gap and 2-Wasserstein distortion criteria; the headline result is a phase-transition partition of the regime axis governed by the dimensionless ratio $\rho = \Sigma_0\, m_0/\sigma^2$. Frøseth (2026h) provides the W_2-geometric refinement on the unrestricted schedule class via the rigid-motion characterisation $\tau^*(x) = \lambda x^2$.

The heterogeneous-returns extension of these geometric criterion families is left for future work. The present paper's decomposition results (Sections 5–6.4) are likely to port into the W_2 framework with the same qualitative structure — the dormant-wealth channel and the differential-incidence proposition are properties of the post-tax SDE, not of the criterion family — but the quantitative substitutability calculation is criterion-specific and will require its own treatment. We do not pursue the W_2 refinement here.

7.6 Position summary

The contribution of this paper to the corpus is therefore twofold. First, it extends the homogeneous-returns design framework of Frøseth (2026f) to type-stratified return primitives, absorbing what would otherwise have been three separate papers (Frøseth (2026c)-style heterogeneous neutrality extension, Frøseth (2026g)-style heterogeneous drift design, and the synthesis) into a single Section 3. Second, it makes precise the empirical limits of Frøseth (2026f)'s homogeneous-returns substitutability claim: under realistic Norwegian asset-class composition at the top of the wealth distribution, with the *fritaksmetoden* channel correcting $\theta(z_H)$ down to the 0.1–0.2 range, the substitutability fails not at a marginal quantitative correction but at the order-of-magnitude level. The wealth tax has a structurally distinct role under heterogeneity that flow taxes cannot replicate, and the framework pins down both the qualitative mechanism and the empirical magnitudes.

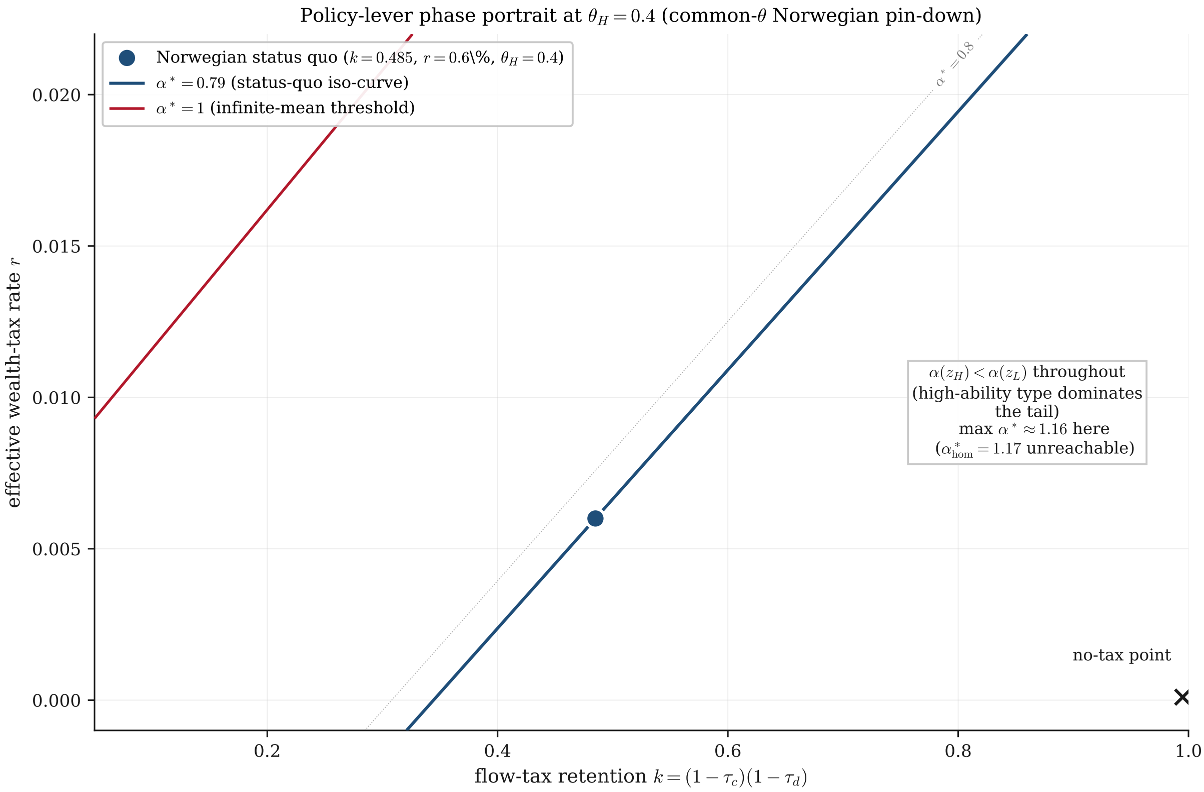

Figure 3: Policy-lever phase portrait under heterogeneous returns — the heterogeneous-returns analogue of Frøseth (2026b, Fig. 2). Plotted at the common-θ Norwegian pin-down $\theta_H = 0.4$. Throughout the empirically plausible (k, r) range shown, the high-ability type carries the fatter tail ($\alpha(z_H) < \alpha(z_L)$) and therefore drives the population dominant Pareto exponent $\alpha^* = \alpha(z_H)$; the type-dominance boundary $\alpha(z_L) = \alpha(z_H)$ lies outside the visible range (it would require an extreme tax structure that suppresses the high-ability type's drift below the lower-ability type's, well off the empirical plate). Two highlighted iso-curves carve the plane into the substantive policy regions: the infinite-mean threshold $\alpha^* = 1$ (red solid) below which the population mean diverges, and the status-quo iso-curve $\alpha^* \approx 0.79$ (blue solid) along which the Pareto exponent matches the Norwegian status quo. The Norwegian policy point $(k, r) = (0.485, 0.6\%)$ (blue disc) sits on the status-quo iso-curve, in the infinite-mean region. The homogeneous-baseline level $\alpha^*_{\text{hom}} \approx 1.17$ that would obtain at $\theta_H = 1$ is unreachable anywhere in this slice — the maximum α^* in the visible (k, r) range is about 0.8 — a substantive finding rather than a plotting artefact: heterogeneity at $\theta_H = 0.4$ shifts the entire α^* surface downward enough that no policy combination (k, r) in the empirically plausible range can recover the homogeneous-baseline tail thickness. Faint dotted grey contours mark intermediate α^* levels. Calibration matches Figure 1.

8 Discussion

The heterogeneous-returns framework of Section 2 treats the ability variable z as a primitive and computes the inequality-criterion consequences of policy changes through it. This treatment is silent on the substantive interpretive question: what is z? The answer matters for policy because the welfare reading of the cross-type reallocation generated by a flow-to-stock substitution (Sections 4–6.5) depends on it. We discuss this here, lean on the developed treatment in Frøseth (2026c, §6.4–6.5), and note the implications that are specific to the redistributive-criterion reading of the present paper.

8.1 Skill versus structural advantage as the source of return heterogeneity

Frøseth (2026c, §6.4) decomposes observed return heterogeneity into three categories.

Category I: Skill in private markets. Persistent return heterogeneity in private markets arises in part from genuine differences in productive ability — judgement in project selection, superior management, deeper sector expertise. The Guvenen et al. (2023) efficiency reading rests on this interpretation: the wealth tax accelerates the reallocation of capital from less productive to more productive hands, and the welfare gain is the productivity gain. Guvenen et al. (2024) sharpen this with an innovation/effort channel: the wealth tax does not tax the upside of productive investment (entrepreneurs keep the excess return), so it incentivises higher innovation effort and entrepreneurial intensity beyond the reallocation effect alone. Under the productive-skill interpretation, this is the strongest theoretical case for wealth taxation. The Frøseth (2026c, §6.4) observation about this category is that skill generates persistent returns *only because* the market is partially inefficient. In a fully efficient market, even highly skilled investors earn zero excess return. Persistence is therefore a joint product of skill and market structure.

Category II: Structural advantage. Return persistence can also arise from features orthogonal to individual skill: scale economies in capital access, sector exposure, vintage effects, local market structure, relationship capital, regulatory protections. Fagereng et al. (2020)'s individual fixed effects on Norwegian register data partly reflect firm characteristics that are not portable across owners. Under this interpretation the use-it-or-lose-it reallocation channels capital toward *positions* rather than toward *people*, and the efficiency reading does not apply.

Category III: Winner-take-all dynamics. In digital markets, near-zero marginal cost, network effects, data advantages, and switching costs translate small initial advantages into large persistent return differences (Rosen, 1981; Brynjolfsson and McAfee, 2014; Autor et al., 2020; De Loecker et al., 2020). Two founders of identical ability launching identical products months apart can earn vastly different returns. The return to the dominant firm reflects structural incumbency, not ongoing skill.

The three-way decomposition $z_i = z_i^{\text{skill}} + z_i^{\text{structure}} + z_i^{\text{platform}}$ (Frøseth, 2026c, eq. (z_decomp)) is pedagogically useful but empirically the components are non-separable: skill operates through market structure (Frøseth, 2026c, eq. (z_interaction)). The implication for the present paper is that the redistributive quantities computed in Sections 4–6.4 do not themselves depend on the decomposition — the joint Pareto formula (17), the cross-type drift gap (19), the substitutability table 5 — but the *policy reading* of those quantities does.

8.2 What the framework pins down and what it does not

The redistributive criteria evaluated in this paper (Gini and top share, plus the population dominant Pareto exponent α^*) are silent on the welfare interpretation of a cross-type reallocation. They pin down the *magnitude* of the reallocation generated by a given tax substitution: holding α^* fixed, raising k shifts mass from the low-ability marginal to the high-ability marginal in iso-criterion units. They do not pin down whether the shift represents a productivity-driven welfare gain (Guvenen et al., 2023's reading), a competitive-distortion welfare loss (the Category-II / Category-III reading), or a mixture.

This silence is policy-relevant. A debate over whether the Norwegian wealth tax is welfare-improving on the use-it-or-lose-it channel cannot be settled by computing α^* or the substitutability table; it requires a prior on the decomposition of observed return heterogeneity in the empirically dominant Norwegian asset classes. The present framework supplies the magnitudes the welfare debate then weighs.

8.3 Implications for the substitutability finding

The §6 substitutability finding of the present paper — that the wealth tax is not substitutable through flow taxation alone at empirically plausible Norwegian top-of- distribution asset-class composition — is robust to the decomposition question, but its policy reading is not.

If skill dominates. If observed top-end return heterogeneity reflects primarily Category I (skill amplified by partial market inefficiency), the substitutability failure *strengthens* the case for the wealth tax: the wealth tax is the instrument that preserves the cross-type drift gap on which the use-it-or-lose-it efficiency channel operates, and flow taxes cannot replicate this even at the calibrated ~ 18 pp dividend-tax compensation level. The infeasibility of substitutability therefore does not argue for removing the wealth tax; it argues that the wealth tax captures a structural lever the flow tax cannot reach, which is welfare-improving when the dominant source of persistence is skill.

If structural advantage dominates. If the dominant source is Category II or III (structural advantage, incumbency, platform rents), the substitutability failure *weakens* the case for the wealth tax in the welfare sense, but does not eliminate it: the wealth tax still reaches the dormant-wealth base that flow taxes structurally cannot reach, and the alternative is not a flow-tax substitute (which is infeasible at the calibrated $\theta(z_H)$) but *either* retaining the wealth tax *or* accepting a policy gap. Under the structural-advantage reading, the wealth tax does not help reallocate capital toward productive holders — it accelerates concentration toward incumbent positions. The redistributive-criterion reading remains favourable (the tax compresses the upper tail) but the welfare reading reverses sign relative to Category I.

The mixed case. The empirically realistic case is a Category-I/II/III mixture, with the share parameters likely sector-specific. Frøseth (2026c, §6.5) argues that the optimal-policy implication is sector-differentiated rates or a complementary combination of wealth taxation (to exploit the skill channel) and competition policy or excess-profit taxes (to address structural rents). The present paper's contribution to that broader argument is the substitutability finding: the homogeneous-returns claim of Frøseth (2026f, §6.6) that the wealth tax is quantitatively replaceable by flow taxes within the historical Norwegian range fails at the order-of-magnitude level once the empirical dormant-wealth share is admitted, *regardless of which decomposition the welfare reading adopts*. The finding is therefore robust to the skill-vs-structure debate.

8.4 Productive control versus passive holding

A related but distinct distinction is documented in the Norwegian emigration evidence of Frøseth (2026a). Wealth ownership decomposes into *active control* (voting rights, strategic oversight, productivity haircut $\eta_A > 0$ on emigration) and *passive holding* (economic rights without operational role, productivity haircut $\eta_P \approx 0$). In the typical Norwegian family firm with dual share classes, the heir emigrates with the B-shares (passive economic ownership) while the controlling parent retains the A-shares (active control); the wealth moves but the productive control does not. The empirical finding of Frøseth (2026a, §5.2) is that most Norwegian top-end emigrating wealth is passive, so the GDP effect of emigration is negligible regardless of total emigrated wealth. For the present paper, the parallel is direct: the type variable z in the productive-skill reading (Category I above) corresponds to active control; the return persistence of passive holders, by contrast, is dominantly Category-II or III, reflecting the underlying firm or platform position rather than the holder's skill. The Frøseth (2026a) active–passive split therefore provides an institutional analogue of the skill-vs-structure split, and the empirical evidence on the passive share at the Norwegian top is informative about the likely dominance of structural sources in the heterogeneous-returns decomposition.

8.5 Open directions

Three extensions of the present paper are most directly suggested by the analysis.

Empirical decomposition of return persistence. The Abowd et al. (1999) person-firm decomposition applied to Norwegian register data could in principle pin down the share of observed return heterogeneity attributable to portable individual ability versus non-portable firm characteristics. This would translate directly into a prior on the welfare sign of the substitutability finding here. Frøseth (2026c, §6.4) sketches the identification strategy; the empirical implementation is a separate research programme.

Multi-bracket and continuous-type extensions. The two-type calibration of Sections 5 and 6 is a stylised summary of a continuous gradient. The natural extension is to a continuous type distribution $F(z)$ on $z \in Z$, with type-conditional asset-class composition $\theta(z)$ as a smooth function. The substitutability calculation then becomes a functional minimisation; whether the substitutability-failure threshold survives a continuous treatment, and where on the type distribution it bites, is an open quantitative question.

Behavioural responses to the schedule. The framework treats the type variable z and the asset-class composition $\theta(z)$ as exogenous to the schedule. In practice both respond to the tax structure: high-ability investors may sort into asset classes that minimise their effective tax burden (a $\theta(z)$ response), and the underlying distribution of z may itself depend on the incentive-to-invest. Frøseth (2026d) characterises the portfolio-choice channel under (C1)–(C3) and shows that proportional wealth taxes are portfolio-neutral in the homogeneous-returns case; whether the analogous neutrality holds under heterogeneity, and whether the *fritaksmetoden*-mediated $\delta_i \approx 0$ row of Table 2 survives an endogenous ownership-structure response, is open.

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

This appendix collects the proofs of the lemmas, theorem, and propositions stated in Sections 3 and 4. Statements are reproduced for reference; reader-facing notation matches the main text.

A.1 Proof of Lemma 1 (Heterogeneous generalised neutrality)

Proof. The strategy is to apply Frøseth (2026d, Theorem 1) type-by-type, verifying that no step of the homogeneous-returns proof relies on across-type assumptions.

Fix a type $z \in Z$. The type-z investor faces the tax-blind asset menu $\{i = 0, 1, \ldots, N\}$ where asset 0 is the risk-free asset with return r_f and assets $1, \ldots, N$ are risky with type-conditional gross returns $\mu_i(z)$ and a covariance matrix $\Sigma(z) = (\sigma_{ij}(z))$. The investor chooses portfolio weights $\boldsymbol{w}(z) = (w_1(z), \ldots, w_N(z))$ to maximise

$$\mathbb{E}\left[\frac{W_T(z)^{1-\gamma}}{1-\gamma}\right], \qquad \gamma > 0,\ \gamma \neq 1,$$

(or its $\gamma = 1$ logarithmic limit) where $W_T(z)$ is terminal wealth at horizon T under the tax system.

Step 1: Excess-return form under (C1)–(C3). By hypothesis (C1)–(C3), the after-tax instantaneous return on risky asset i for the type-z investor is

$$R_i^{\text{after}}(z) \;=\; k\,\mu_i(z) - \tau_w + \xi_i,$$

where $k = (1-\tau_c)(1-\tau_d)$ is the corporate–dividend retention factor of (C1), $-\tau_w$ is the proportional wealth-tax drag of (C2), and ξ_i is a zero-mean shock with covariance $k\,\sigma_{ij}(z)$ (the symmetric loss-offset condition (C3) ensures the variance also scales by k). The risk-free return after tax becomes $k\,r_f - \tau_w$ (treating r_f as taxable interest). The after-tax excess-return vector relative to the risk-free asset is therefore

$$\boldsymbol{R}^{\mathrm{ex}}(z) \;=\; \boldsymbol{R}^{\mathrm{after}}(z) - (k\,r_f - \tau_w)\,\mathbf{1} \;=\; k\,\big(\boldsymbol{\mu}(z) - r_f\,\mathbf{1}\big),$$

i.e. the wealth-tax drag $-\tau_w$ cancels uniformly across assets and the residual excess-return vector is k times the pre-tax excess-return vector. The covariance matrix of $\boldsymbol{R}^{\mathrm{after}}(z)$ is $k\,\Sigma(z)$.

Step 2: Mean–variance optimum. For CRRA utility, in either the small-noise / continuous-time limit or for log utility ($\gamma = 1$) exactly, the optimal portfolio weights satisfy

$$\boldsymbol{w}^*(z) \;\propto\; \big[k\,\Sigma(z)\big]^{-1}\boldsymbol{R}^{\mathrm{ex}}(z) \;=\; \big[k\,\Sigma(z)\big]^{-1} k\,\big(\boldsymbol{\mu}(z) - r_f\,\mathbf{1}\big) \;=\; \Sigma(z)^{-1}\,\big(\boldsymbol{\mu}(z) - r_f\,\mathbf{1}\big),$$

which is independent of k and identical to the no-tax ($k = 1$, $\tau_w = 0$) optimum. For general γ, the proportionality factor depends on γ but the *direction* of $\boldsymbol{w}^*(z)$ in $\mathbb{R}^N$ is the same. The wealth-tax rate τ_w enters neither the direction nor the magnitude (it shifts the risk-free rate uniformly).

Step 3: Type-by-type independence. The argument in Steps 1–2 uses only the type-z primitives $(\mu_i(z), \Sigma(z))$ and the tax parameters (τ_c, τ_d, τ_w); no across-type assumption is invoked. Hence $\boldsymbol{w}^*(z)$ is independent of (τ_c, τ_d, τ_w) for each z separately, with the optimum determined by the no-tax solution at the type-z primitives.

Step 4: Consequent post-tax SDE on log-wealth. Under the optimal weights, the type-z investor's wealth satisfies

$$\frac{\mathrm{d}W_t(z)}{W_t(z)} \;=\; \big[k\,\mu(z) - \tau_w\big]\,\mathrm{d}t + \sqrt{k}\,\sigma(z)\,\mathrm{d}B_t^{(z)},$$

where $\mu(z) = \boldsymbol{w}^*(z)'\boldsymbol{\mu}(z)$ is the optimal-portfolio mean return and $\sigma(z) = \sqrt{\boldsymbol{w}^*(z)'\Sigma(z)\boldsymbol{w}^*(z)}$ is the optimal-portfolio volatility. By Itô's lemma applied to $Y_t = \log W_t(z)$,

$$\mathrm{d}Y_t \;=\; \big[k\,\mu(z) - \tau_w - \tfrac{1}{2}\,k\,\sigma(z)^2\big]\,\mathrm{d}t + \sqrt{k}\,\sigma(z)\,\mathrm{d}B_t^{(z)} \;=\; \big[k\,m_0(z) - \tau_w\big]\,\mathrm{d}t + \sqrt{k}\,\sigma(z)\,\mathrm{d}B_t^{(z)},$$

recovering (4) after generalising the constant rate τ_w to the bracket schedule $r(e^{Y_t}) = r(W_t)$ (which is type-blind by (C2)). $\square$

A.2 Proof of Lemma 2 (Implementability-constrained design)

Proof. Step 1: Joint stationary density. Under Assumption 1 (the type z is private and the tax authority observes only wealth X_t, so the schedule is a function $r : (0,\infty) \to [0, r_{\max}]$ of wealth alone), each type's log-wealth process satisfies (4). Augmenting with a type-conditional demographic source–sink at rate $\delta(z) > 0$ — births at a fixed log-wealth y_0, deaths uniformly across the distribution — the joint stationary density $\pi(y, z)$ on $\mathbb{R} \times Z$ satisfies the joint Fokker–Planck equation

$$\partial_y\big[(v_{\mathrm{eff}}(z) - r(e^y))\,\pi(y,z)\big] + \delta(z)\,\pi(y,z) \;=\; D_{\mathrm{eff}}(z)\,\partial_y^2\pi(y,z) + \delta(z)\,F(\mathrm{d}z)\cdot\delta_{y_0}(y),$$

where $v_{\text{eff}}(z) = k\, m_0(z) - r$ and $D_{\text{eff}}(z) = k\, \sigma(z)^2/2$ are the type-conditional drift and diffusion of (16). Existence and uniqueness of $\pi(y, z)$ under standard regularity conditions (sufficient confinement from $\delta > 0$, locally bounded r) follows from Frøseth (2026f, §6.5)'s argument applied type-by-type.

Step 2: Marginal density and type-conditioning. Integrating in z against $F(\mathrm{d}z)$ gives the marginal log-wealth density

$$\pi(y) \;=\; \int_Z F(\mathrm{d}z)\, \pi(y \mid z) \;=\; \int_Z F(\mathrm{d}z)\, \frac{\pi(y, z)}{F(\mathrm{d}z)}.$$

The type-conditional density $\pi(y \mid z) = \pi(y, z)/F(\mathrm{d}z)$ captures the conditional distribution of log-wealth given type. The type-conditioning moments

$$\mathbb{E}[m_0(z) \mid Y = y] \;=\; \frac{\int_Z F(\mathrm{d}z)\, m_0(z)\, \pi(y \mid z)}{\pi(y)}, \qquad \mathbb{E}[\sigma(z)^2 \mid Y = y] \;=\; \frac{\int_Z F(\mathrm{d}z)\, \sigma(z)^2\, \pi(y \mid z)}{\pi(y)},$$

are functionals of the schedule $r(\cdot)$ because $\pi(y \mid z)$ depends on r through the type-z stationary density.

Step 3: Reduction to marginal extremisation. The inequality criteria $\Phi[\pi]$ (Gini, top share, population dominant Pareto exponent) depend on the marginal density $\pi(y)$ alone, not on the joint $\pi(y, z)$. Hence the design problem is

$$r^* \;=\; \arg\min_{r \in \mathcal{R}} \Phi[\pi(\cdot\,; r)] \quad \text{subject to revenue and other constraints,}$$

where $\mathcal{R}$ is the admissible schedule class and $\pi(\cdot\,; r)$ is the marginal density induced by r via the joint FP system.

Step 4: Fixed-point structure of the optimality condition. The first-order condition for the extremisation $\Phi[\pi(\cdot\,; r)]$ (computed via functional derivatives w.r.t. $r(\cdot)$) takes the form $\delta\Phi/\delta r\,|_{r=r^*} = 0$, which after the chain rule through $\pi(\cdot\,; r^*)$ reduces to a stationarity condition involving the type-conditioning moments $\mathbb{E}[m_0(z) \mid Y = y]$ and $\mathbb{E}[\sigma(z)^2 \mid Y = y]$ evaluated at $\pi(\cdot\,; r^*)$. Since these moments themselves depend on r^* via $\pi(\cdot\,; r^*)$, the optimality condition is implicit in r^*: the optimal schedule must satisfy the FOC at the steady state it generates. This is the fixed-point structure asserted in the lemma. □

A.3 Proof of Lemma 3 (Heterogeneous neutrality with dormant wealth)

Proof. Step 1: Sub-portfolio decomposition. Decompose the type-z investor's wealth into income-producing and dormant sub-portfolios with shares $\theta(z) \in [0, 1]$ and $1 - \theta(z)$ respectively. The dynamics of each sub-portfolio are characterised by separate effective retention factors:

- Income-producing sub-portfolio: faces the (C1)–(C3) flow-tax envelope of Lemma 1, with drift coefficient $k\, \mu(z)$ and diffusion variance $k\, \sigma(z)^2$ on the underlying wealth (before Itô correction).
- Dormant sub-portfolio: faces no flow tax (the income-flow ratio of the dormant assets is

identically zero by definition), with drift coefficient $\mu(z)$ and diffusion variance $\sigma(z)^2$.

The wealth tax acts uniformly on raw wealth at rate r for both sub-portfolios; this is condition (C3) extended to the dormant base, an assumption made explicit in the lemma's hypothesis (the wealth tax assesses raw wealth uniformly across asset classes).

Step 2: Aggregation under the share-weighted-variance convention. The combined log-wealth dynamics are obtained by share-weighted aggregation of the sub-portfolio drifts and variances. For the drift on raw wealth:

$$\theta(z)\cdot k\,\mu(z) + (1-\theta(z))\cdot\mu(z) \;=\; \mu(z)\,[1-\theta(z)\,(1-k)] \;=\; k_{\text{eff}}(z)\,\mu(z),$$

recovering $k_{\text{eff}}(z) = 1-\theta(z)(1-k)$ as defined in (12). For the diffusion variance on log-wealth, under the share-weighted-variance convention common in the wealth-dynamics literature (Frøseth, 2026f, §6.5 for the homogeneous analogue):

$$\theta(z)\cdot k\,\sigma(z)^2 + (1-\theta(z))\cdot\sigma(z)^2 \;=\; k_{\text{eff}}(z)\,\sigma(z)^2.$$

The corresponding diffusion coefficient on log-wealth is $\sqrt{k_{\text{eff}}(z)}\,\sigma(z)$.

Step 3: Itô correction and reassembly. On the raw wealth level, the share-weighted dynamics are

$$\frac{\mathrm{d}W_t(z)}{W_t(z)} \;=\; [k_{\text{eff}}(z)\,\mu(z) - r(e^{Y_t})]\,\mathrm{d}t + \sqrt{k_{\text{eff}}(z)}\,\sigma(z)\,\mathrm{d}W_t^{(z)}.$$

Applying Itô's lemma to $Y_t = \log W_t(z)$ with the diffusion coefficient $\sqrt{k_{\text{eff}}(z)}\,\sigma(z)$:

$$\begin{aligned}
\mathrm{d}Y_t &= [k_{\text{eff}}(z)\,\mu(z) - r(e^{Y_t}) - \tfrac{1}{2}\,k_{\text{eff}}(z)\,\sigma(z)^2]\,\mathrm{d}t + \sqrt{k_{\text{eff}}(z)}\,\sigma(z)\,\mathrm{d}W_t^{(z)}\\
&= [k_{\text{eff}}(z)\,(\mu(z) - \tfrac{\sigma(z)^2}{2}) - r(e^{Y_t})]\,\mathrm{d}t + \sqrt{k_{\text{eff}}(z)}\,\sigma(z)\,\mathrm{d}W_t^{(z)}\\
&= [k_{\text{eff}}(z)\,m_0(z) - r(e^{Y_t})]\,\mathrm{d}t + \sqrt{k_{\text{eff}}(z)}\,\sigma(z)\,\mathrm{d}W_t^{(z)},
\end{aligned}$$

which is (13).

Step 4: Limiting cases. At $\theta(z) = 1$ (fully income-producing portfolio), $k_{\text{eff}}(z) = k$ and the SDE reduces to the heterogeneous-neutrality SDE (4) of Lemma 1. At $\theta(z) = 0$ (fully dormant), $k_{\text{eff}}(z) = 1$ and the SDE reduces to the no-flow-tax dynamics. Both limits are consistent with the share-weighted-variance convention. □

A.4 Proof of Theorem 1 (Type-conditional joint Pareto exponent)

Proof. Step 1: Type-conditional steady-state ODE. Under Lemma 1 and a type-conditional demographic source–sink at rate $\delta(z) > 0$ at log-wealth y_0, the type-conditional stationary density $\pi(y \mid z)$ satisfies the steady-state Fokker–Planck equation

$$-\partial_y[v_{\text{eff}}(z)\,\pi(y\mid z)] + D_{\text{eff}}(z)\,\partial_y^2\,\pi(y\mid z) + \delta(z)\,\delta_{y_0}(y) - \delta(z)\,\pi(y\mid z) \;=\; 0,$$

where the first two terms are the drift–diffusion generator of the log-wealth process under a flat bracket $r(\cdot) = r$ (so $v_{\text{eff}}(z)$ is constant in y) and the latter two terms model demographic births at y_0 and uniform-rate deaths.

Step 2: Upper-tail homogeneous-coefficient ODE. For $y > y_0$, the source $\delta(z)\,\delta_{y_0}(y)$ vanishes and the equation reduces to the homogeneous-coefficient second-order linear ODE

$$D_{\text{eff}}(z)\,\pi''(y \mid z) - v_{\text{eff}}(z)\,\pi'(y \mid z) - \delta(z)\,\pi(y \mid z) \;=\; 0. \tag{24}$$

Step 3: Exponential ansatz and characteristic equation. Substituting $\pi(y \mid z) = A(z)\,e^{-\alpha(z)\,y}$ into (24):

$$D_{\text{eff}}(z)\,\alpha(z)^2\,A(z)\,e^{-\alpha(z)y} + v_{\text{eff}}(z)\,\alpha(z)\,A(z)\,e^{-\alpha(z)y} - \delta(z)\,A(z)\,e^{-\alpha(z)y} \;=\; 0.$$

Dividing through by $A(z)\,e^{-\alpha(z)y}$ yields the characteristic equation

$$D_{\text{eff}}(z)\,\alpha(z)^2 + v_{\text{eff}}(z)\,\alpha(z) - \delta(z) \;=\; 0. \tag{25}$$

Step 4: Selection of the positive root. The quadratic (25) has two roots:

$$\alpha(z) \;=\; \frac{-v_{\text{eff}}(z) \pm \sqrt{v_{\text{eff}}(z)^2 + 4\,D_{\text{eff}}(z)\,\delta(z)}}{2\,D_{\text{eff}}(z)}.$$

For $\pi(y \mid z)$ to be a normalisable density supported on $y > y_0$ (decaying as $y \to \infty$), $\alpha(z)$ must be positive. The discriminant $v_{\text{eff}}(z)^2 + 4\,D_{\text{eff}}(z)\,\delta(z) > 0$ strictly (since $D_{\text{eff}}(z), \delta(z) > 0$), so $\sqrt{\cdot} > |v_{\text{eff}}(z)|$ strictly, and the positive root is the one with the $+$ sign:

$$\alpha(z) \;=\; \frac{-v_{\text{eff}}(z) + \sqrt{v_{\text{eff}}(z)^2 + 4\,D_{\text{eff}}(z)\,\delta(z)}}{2\,D_{\text{eff}}(z)} \;>\; 0,$$

which is (17). The negative root is rejected because the corresponding density would grow exponentially in y.

Step 5: Boundary matching. The constant $A(z)$ is determined by matching $\pi(y \mid z)$ at $y = y_0$ to the left-tail solution (which involves the second linearly independent root of (25), the negative one). This determines $A(z)$ but does not affect the upper-tail exponent $\alpha(z)$, which is the object of the theorem. □

A.5 Proof of Corollary 1 (Homogeneous recovery)

Proof. At a single representative type $z \equiv z_0$ with primitives $(m_0(z_0), \sigma(z_0)) = (m_0, \sigma)$ and demographic rate $\delta(z_0) = \delta$, formula (17) reduces verbatim to

$$\alpha(z_0) \;=\; \frac{-(k\,m_0 - r) + \sqrt{(k\,m_0 - r)^2 + 2\,k\,\sigma^2\,\delta}}{k\,\sigma^2},$$

which is the joint Pareto formula of Frøseth (2026f, §6.5) under the homogeneous-returns assumption. The numerical check at the calibrated Norwegian parameters ($m_0 = 0.06$, $\sigma = 0.28$, $k = 0.485$, $r = 0.006$, $\delta = 0.062$):

$$\begin{aligned}
v_{\text{eff}} &= 0.485 \cdot 0.06 - 0.006 = 0.0231, \\
D_{\text{eff}} &= 0.485 \cdot 0.0784/2 = 0.01901, \\
\alpha(z_0) &= \frac{-0.0231 + \sqrt{0.0231^2 + 4 \cdot 0.01901 \cdot 0.062}}{2 \cdot 0.01901} = \frac{-0.0231 + 0.07244}{0.03802} = 1.298,
\end{aligned}$$

matching the Frøseth (2026f, §6.6) status-quo value of 1.30 (rounded to two decimals). □

A.6 Proof of Corollary 2 (Bracket-strand recovery)

Proof. Take $\delta(z) \to 0$ in (17). The discriminant reduces to $v_{\text{eff}}(z)^2$ in this limit, so

$$\alpha(z) \;\to\; \frac{-v_{\text{eff}}(z) + |v_{\text{eff}}(z)|}{2\, D_{\text{eff}}(z)}.$$

Two regimes by the sign of $v_{\text{eff}}(z) = k\, m_0(z) - r$:

Case 1 (finite-mean regime, $r > k\, m_0(z)$): $v_{\text{eff}}(z) < 0$, so $|v_{\text{eff}}(z)| = -v_{\text{eff}}(z) = r - k\, m_0(z)$, and

$$\alpha(z) \;=\; \frac{-v_{\text{eff}}(z) + |v_{\text{eff}}(z)|}{2\, D_{\text{eff}}(z)} \;=\; \frac{-2\, v_{\text{eff}}(z)}{2\, D_{\text{eff}}(z)} \;=\; \frac{r - k\, m_0(z)}{k\, \sigma(z)^2/2} \;=\; \frac{2\,(r - k\, m_0(z))}{k\, \sigma(z)^2} \;=\; \frac{2\,(r/k - m_0(z))}{\sigma(z)^2},$$

which is the bracket-strand exponent $\alpha_2(z)$ of Frøseth (2026f, Definition 3) evaluated at type-z primitives.

Case 2 (infinite-mean regime, $r < k\, m_0(z)$): $v_{\text{eff}}(z) > 0$, so $|v_{\text{eff}}(z)| = v_{\text{eff}}(z)$, and

$$\alpha(z) \;=\; \frac{-v_{\text{eff}}(z) + v_{\text{eff}}(z)}{2\, D_{\text{eff}}(z)} \;=\; 0,$$

the degenerate flat-tail limit. In this regime the bracket-strand formula does not apply (the upper tail has infinite mean and the demographic confinement $\delta > 0$ is needed for a well-defined Pareto exponent).

In either case the corollary's claim — the $\delta(z) \to 0$ recovery of the bracket-strand exponent under the finite-mean condition — holds. □

A.7 Proof of Proposition 1 (Population dominant Pareto exponent)

Proof. Step 1: Type-conditional asymptotics. By Theorem 1, for each $z \in \operatorname{supp}(F)$,

$$\pi(y \mid z) \;=\; A(z)\, e^{-\alpha(z)\, y}\, (1 + o(1)) \quad \text{as} \quad y \to \infty,$$

where the $o(1)$ correction captures the lower-order corrections to the bracket-strand exponent (subleading in y).

Step 2: Marginal density. The population marginal density is the type average

$$\pi(y) \;=\; \int_Z F(\mathrm{d}z)\, \pi(y \mid z).$$

Substituting the type-conditional asymptotics:

$$\begin{aligned}\pi(y) &= \int_Z F(\mathrm{d}z)\, A(z)\, e^{-\alpha(z)\, y}\, (1 + o(1)) \\ &= e^{-\alpha^*\, y} \int_Z F(\mathrm{d}z)\, A(z)\, e^{-(\alpha(z) - \alpha^*)\, y}\, (1 + o(1)),\end{aligned}$$

factoring out $e^{-\alpha^*\, y}$ where $\alpha^* = \inf_{z \in \mathrm{supp}(F)} \alpha(z) = \alpha(z^*)$ for the inf-attaining type z^*.

Step 3: Asymptotic behaviour of the integral. The exponent $\alpha(z) - \alpha^* \geq 0$ for every $z \in \mathrm{supp}(F)$, with equality only at z^* (or the inf-attaining set, treated below). Distinguish two cases by the structure of F near z^*:

Case 1 (atomic F at z^*): If $F(\{z^*\}) > 0$, the mass at z^* contributes a constant $F(\{z^*\}) \cdot A(z^*)$ to the integral as $y \to \infty$ (since the integrand at z^* equals $A(z^*)$ and away from z^* the exponent $\alpha(z) - \alpha^* > 0$ kills the contribution). The remainder $\int_{\{z \neq z^*\}} F(\mathrm{d}z)\, A(z)\, e^{-(\alpha(z) - \alpha^*)\, y}$ is bounded by $\sup A \cdot e^{-\eta\, y}$ for some $\eta > 0$ (any positive lower bound on $\alpha(z) - \alpha^*$ off the atom at z^*), so the remainder is exponentially smaller than the atom contribution.

Case 2 (continuous F attaining its inf at z^*): For z in a neighbourhood of z^*, expand $\alpha(z) - \alpha^*$ in a Taylor series around z^*. If z^* is an interior point of $\mathrm{supp}(F)$ at which α has a non-degenerate minimum (positive-definite Hessian), $\alpha(z) - \alpha^* \approx \frac{1}{2}\, (z - z^*)' H (z - z^*)$ for the Hessian H of α at z^*. Laplace's method then gives

$$\int_Z F(\mathrm{d}z)\, A(z)\, e^{-(\alpha(z) - \alpha^*)\, y} \;\sim\; \frac{(2\pi/y)^{d/2}}{\sqrt{|H|}}\, f_F(z^*)\, A(z^*) \qquad (y \to \infty),$$

where d is the dimension of Z and f_F is the density of F at z^*. The y-decay of this integral is power-law ($y^{-d/2}$), much slower than the exponential $e^{-\alpha^*\, y}$ in the prefactor.

Step 4: Combined asymptotics. In Case 1,

$$\pi(y) \;\sim\; F(\{z^*\}) \cdot A(z^*) \cdot e^{-\alpha^*\, y}.$$

In Case 2,

$$\pi(y) \;\sim\; C \cdot y^{-d/2} \cdot e^{-\alpha^*\, y}, \qquad C = (2\pi)^{d/2} / \sqrt{|H|} \cdot f_F(z^*) \cdot A(z^*).$$

In both cases $\log \pi(y) = -\alpha^*\, y + O(\log y)$, so the Pareto exponent of the population marginal is exactly $\alpha^* = \min_z \alpha(z) = \alpha(z^*)$. The prefactor A^* is given by the case-specific expression in the proposition statement. $\square$

A.8 Proof of Proposition 2 (Differential incidence in the cross-type drift gap)

Proof. The cross-type drift gap is $\Delta v = v_{\text{eff}}(z_H) - v_{\text{eff}}(z_L) = [k\, m_0(z_H) - r] - [k\, m_0(z_L) - r] = k\,[m_0(z_H) - m_0(z_L)]$. Differentiating in r:

$$\frac{\partial \Delta v}{\partial r} \;=\; \frac{\partial}{\partial r}\big(k\,[m_0(z_H) - m_0(z_L)]\big) \;=\; 0,$$

since r does not appear (the wealth-tax rate cancels in the cross-type difference). Differentiating in k:

$$\frac{\partial \Delta v}{\partial k} \;=\; m_0(z_H) - m_0(z_L) \;=\; \mu(z_H) - \mu(z_L) \;>\; 0$$

under common σ (so that $m_0(z) - m_0(z') = \mu(z) - \mu(z')$) and the maintained hypothesis $\mu(z_H) > \mu(z_L)$. □

A.9 Proof of Proposition 3 (α-gap response)

Proof. We work under the maintained hypotheses common volatility $\sigma(z) \equiv \sigma$, common turnover rate $\delta(z) \equiv \delta > 0$, and $m_0(z_L) < m_0(z_H)$ (equivalently $v_{\text{eff}}(z_L) < v_{\text{eff}}(z_H)$ at any common (k, r)).

Step 1: Implicit derivative formula. Differentiate the characteristic equation (25) implicitly with respect to a parameter $\lambda \in \{r, k\}$:

$$\frac{\partial D_{\text{eff}}}{\partial \lambda}\,\alpha^2 + 2\,D_{\text{eff}}\,\alpha\,\frac{\partial \alpha}{\partial \lambda} + \frac{\partial v_{\text{eff}}}{\partial \lambda}\,\alpha + v_{\text{eff}}\,\frac{\partial \alpha}{\partial \lambda} \;=\; 0,$$

which solves to

$$\frac{\partial \alpha}{\partial \lambda} \;=\; -\,\frac{(\partial D_{\text{eff}}/\partial \lambda)\,\alpha^2 + (\partial v_{\text{eff}}/\partial \lambda)\,\alpha}{2\,D_{\text{eff}}\,\alpha + v_{\text{eff}}} \;=\; -\,\frac{(\partial D_{\text{eff}}/\partial \lambda)\,\alpha^2 + (\partial v_{\text{eff}}/\partial \lambda)\,\alpha}{S}, \tag{26}$$

using $S = 2\,D_{\text{eff}}\,\alpha + v_{\text{eff}} = \sqrt{v_{\text{eff}}^2 + 4\,D_{\text{eff}}\,\delta}$ from the positive-root selection of Theorem 1.

Step 2: Wealth-tax derivative $\partial\alpha/\partial r$. With $\partial v_{\text{eff}}/\partial r = -1$ and $\partial D_{\text{eff}}/\partial r = 0$, (26) gives

$$\frac{\partial \alpha(z)}{\partial r} \;=\; -\,\frac{0 \cdot \alpha(z)^2 - \alpha(z)}{S(z)} \;=\; \frac{\alpha(z)}{S(z)}.$$

Differencing,

$$\frac{\partial \Delta\alpha}{\partial r} \;=\; \frac{\alpha(z_L)}{S(z_L)} - \frac{\alpha(z_H)}{S(z_H)},$$

which is (21).

Step 3: Sign of $\partial\Delta\alpha/\partial r$. Define $f(v) := \alpha(v)/S(v)$ at fixed D_{eff} and δ. Using $\alpha(v) = (-v + S(v))/(2D_{\text{eff}})$ and $S(v) = \sqrt{v^2 + 4D_{\text{eff}}\delta}$:

$$f(v) \;=\; \frac{-v + S(v)}{2\,D_{\text{eff}}\,S(v)} \;=\; \frac{1}{2\,D_{\text{eff}}}\Big(1 - \frac{v}{S(v)}\Big).$$

Since $S(v)$ is increasing in $|v|$, for $v > 0$ we have $v/S(v)$ increasing in v (the ratio $v/\sqrt{v^2 + c}$ is monotonic in v for $c > 0$), so $1 - v/S(v)$ is decreasing in v. For $v < 0$, $v/S(v)$ is also increasing in v (toward zero from below), so $1 - v/S(v)$ is again decreasing in v. In either case $f(v)$ is monotonically decreasing in v. Hence $v_{\text{eff}}(z_L) < v_{\text{eff}}(z_H) \implies f(v_{\text{eff}}(z_L)) > f(v_{\text{eff}}(z_H))$, i.e. $\alpha(z_L)/S(z_L) > \alpha(z_H)/S(z_H)$, and $\partial\Delta\alpha/\partial r > 0$ strictly.

Step 4: Flow-tax-retention derivative $\partial\alpha/\partial k$. With $\partial v_{\text{eff}}/\partial k = m_0(z)$ and $\partial D_{\text{eff}}/\partial k = \sigma^2/2$, (26) gives

$$\frac{\partial\alpha(z)}{\partial k} \;=\; -\,\frac{(\sigma^2/2)\,\alpha(z)^2 + m_0(z)\,\alpha(z)}{S(z)} \;=\; -\,\frac{\alpha(z)}{S(z)}\Big(m_0(z) + \tfrac{\sigma^2}{2}\,\alpha(z)\Big).$$

Differencing,

$$\frac{\partial\Delta\alpha}{\partial k} \;=\; -\,\frac{\alpha(z_L)}{S(z_L)}\Big(m_0(z_L) + \tfrac{\sigma^2}{2}\,\alpha(z_L)\Big) + \frac{\alpha(z_H)}{S(z_H)}\Big(m_0(z_H) + \tfrac{\sigma^2}{2}\,\alpha(z_H)\Big),$$

which is (22). The sign is parameter-dependent (the z_L term has larger α/S but smaller m_0 and larger α than the z_H term, so the products compete); numerical evaluation at the Section 5.1 calibration gives $\partial\Delta\alpha/\partial k \approx -0.17 < 0$.

Step 5: Iso-α^* derivative. Along the curve $\{(k,r) : \alpha(z_H; k, r) = \alpha^*\}$, r is implicitly a function of k, with derivative determined by total differentiation of the constraint:

$$\frac{\partial\alpha(z_H)}{\partial k} + \frac{\partial\alpha(z_H)}{\partial r}\,\frac{\mathrm{d}r}{\mathrm{d}k} \;=\; 0,$$

giving

$$\frac{\mathrm{d}r}{\mathrm{d}k} \;=\; -\,\frac{\partial\alpha(z_H)/\partial k}{\partial\alpha(z_H)/\partial r} \;=\; -\,\frac{-(\alpha(z_H)/S(z_H))(m_0(z_H) + \sigma^2\alpha(z_H)/2)}{\alpha(z_H)/S(z_H)} \;=\; m_0(z_H) + \tfrac{\sigma^2}{2}\,\alpha(z_H).$$

The total derivative of $\Delta\alpha$ along the iso-α^* curve is

$$\begin{aligned}
\frac{\mathrm{d}\Delta\alpha}{\mathrm{d}k}\bigg|_{\alpha^*\text{ fixed}} &= \frac{\partial\Delta\alpha}{\partial k} + \frac{\partial\Delta\alpha}{\partial r}\,\frac{\mathrm{d}r}{\mathrm{d}k} \\
&= \Big[-\,\frac{\alpha(z_L)}{S(z_L)}\,(m_0(z_L) + \tfrac{\sigma^2}{2}\,\alpha(z_L)) + \frac{\alpha(z_H)}{S(z_H)}\,(m_0(z_H) + \tfrac{\sigma^2}{2}\,\alpha(z_H))\Big] \\
&\quad + \Big[\frac{\alpha(z_L)}{S(z_L)} - \frac{\alpha(z_H)}{S(z_H)}\Big]\cdot(m_0(z_H) + \tfrac{\sigma^2}{2}\,\alpha(z_H)).
\end{aligned}$$

The two $\alpha(z_H)/S(z_H)\cdot(m_0(z_H) + \sigma^2\alpha(z_H)/2)$ terms cancel, leaving

$$\begin{aligned}
\frac{\mathrm{d}\Delta\alpha}{\mathrm{d}k}\bigg|_{\alpha^*\text{ fixed}} &= \frac{\alpha(z_L)}{S(z_L)}\,[(m_0(z_H) + \tfrac{\sigma^2}{2}\,\alpha(z_H)) - (m_0(z_L) + \tfrac{\sigma^2}{2}\,\alpha(z_L))] \\
&= \frac{\alpha(z_L)}{S(z_L)}\,[(m_0(z_H) - m_0(z_L)) - \tfrac{\sigma^2}{2}\,(\alpha(z_L) - \alpha(z_H))] \\
&= \frac{\alpha(z_L)}{S(z_L)}\,(\Delta m_0 - \tfrac{\sigma^2}{2}\,\Delta\alpha),
\end{aligned}$$

which is (23).

Step 6: Strict positivity of the iso-α^* derivative for $\delta > 0$. The factor $\alpha(z_L)/S(z_L) > 0$. We must show $\Delta m_0 - (\sigma^2/2)\,\Delta\alpha > 0$. By Corollary 4 (proved below), in the bracket-strand limit $\delta \to 0$ the gap $\Delta\alpha = 2\Delta m_0/\sigma^2$. For any $\delta > 0$, $\Delta\alpha < 2\Delta m_0/\sigma^2$ strictly, because the demographic confinement $\delta > 0$ pushes both $\alpha(z_L)$ and $\alpha(z_H)$ toward a common limit $\sqrt{\delta/D_{\text{eff}}}$ as $\delta \to \infty$, narrowing the gap monotonically:

$$\Delta\alpha \;<\; 2\,\Delta m_0/\sigma^2 \quad\iff\quad \tfrac{\sigma^2}{2}\,\Delta\alpha \;<\; \Delta m_0.$$

Hence $\Delta m_0 - (\sigma^2/2)\,\Delta\alpha > 0$ for any $\delta > 0$, and the iso-α^* derivative is strictly positive.

The closed-form expressions (21), (22), and (23) have been verified symbolically using `sympy`; the verification script is `notes/proof_verification.py` in the source folder, and runs all 15 symbolic checks for the closed-form expressions in this appendix. □

A.10 Proof of Corollary 4 (Bracket-strand recovery and degenerate α-level invariance)

Proof. By Corollary 2, in the bracket-strand limit $\delta \to 0$ with the finite-mean condition $r > k\,m_0(z)$ for all z, the type-conditional Pareto exponent becomes

$$\alpha(z) \;=\; \frac{2\,(r - k\,m_0(z))}{k\,\sigma^2} \;=\; \frac{2\,r}{k\,\sigma^2} - \frac{2\,m_0(z)}{\sigma^2}.$$

The first term $2r/(k\sigma^2)$ is type-blind. The second term $-2m_0(z)/\sigma^2$ is type-dependent and independent of (k, r). Differencing across types:

$$\Delta\alpha \;=\; \alpha(z_L) - \alpha(z_H) \;=\; -\frac{2\,m_0(z_L)}{\sigma^2} + \frac{2\,m_0(z_H)}{\sigma^2} \;=\; \frac{2\,\Delta m_0}{\sigma^2},$$

which is independent of both r and k. Hence $\partial\Delta\alpha/\partial r = \partial\Delta\alpha/\partial k = 0$ in this limit. The iso-α^* derivative (23) also vanishes because the closed-form factor reduces to $\alpha(z_L)/S(z_L) \cdot (\Delta m_0 - \sigma^2/2 \cdot 2\Delta m_0/\sigma^2) = \alpha(z_L)/S(z_L) \cdot 0 = 0$.

The bracket-strand limit therefore recovers the original proposition's null-derivative claims as a degenerate zero-derivative special case, and the joint formula with $\delta > 0$ generates the non-trivial α-level response. □

A.11 Proof of Corollary 3 (Dormant-wealth refinement)

Proof. Under Lemma 3, the post-tax log-wealth SDE for type z has effective drift coefficient $k_{\text{eff}}(z)\,m_0(z)$ and effective diffusion variance $k_{\text{eff}}(z)\,\sigma(z)^2$, with $k_{\text{eff}}(z) = 1 - \theta(z)(1-k)$. Repeating the proof of Theorem 1 with k replaced by $k_{\text{eff}}(z)$ in both v_{eff} and D_{eff} — i.e. $v_{\text{eff}}(z) = k_{\text{eff}}(z)\,m_0(z) - r$ and $D_{\text{eff}}(z) = k_{\text{eff}}(z)\,\sigma(z)^2/2$ — the same characteristic equation (25) holds with the modified definitions, and the positive-root selection gives (17) verbatim. The flow-tax channel enters through $k_{\text{eff}}(z)$ (which depends on the type-conditional dormant share via $\theta(z)$); the wealth-tax channel enters through r (still type-blind); the demographic channel enters through $\delta(z)$ (here common across types). □

References

Rolf Aaberge, Jørgen Modalsli, and Edda Solbakken. Measuring long-run wealth inequality: Empirical results for Norway 1912–2019. Discussion Papers 1028, Statistics Norway, October 2025.

John M. Abowd, Francis Kramarz, and David N. Margolis. High wage workers and high wage firms. *Econometrica*, 67(2):251–333, 1999. doi: 10.1111/1468-0262.00020.

David Autor, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics*, 135(2):645–709, 2020. doi: 10.1093/qje/qjz025.

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.

Petter Bjerksund and Guttorm Schjelderup. Aksjonærmodellen og fritaksmetoden: Et to-hodet troll? *Samfunnsøkonomen*, 135(4):53–72, 2021.

Petter Bjerksund, Arnt Ove Hopland, and Guttorm Schjelderup. The taxation of Norway's richest. *Nordic Tax Journal*, 2024(s1):72–82, 2024. doi: 10.2478/ntaxj-2024-0006.

Erik Brynjolfsson and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton, 2014.

Jan De Loecker, Jan Eeckhout, and Gabriel Unger. The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics*, 135(2):561–644, 2020. doi: 10.1093/qje/qjz041.

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.

Anders G. Frøseth. Tax migration as social contagion: A tipping-point model with application to the Scandinavian wealth tax debate. 2026a. arXiv:2607.03868 [physics.soc-ph].

Anders G. Frøseth. Minimum-distortion wealth taxation, i: Information-theoretic versus transport-geometric optimality on the proportional class. Working paper, 2026b.

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

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

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. From gravity to confinement: Wealth redistribution as optimal drift design in the Fokker–Planck framework. 2026g. arXiv:2607.06153 [physics.soc-ph].

Anders G. Frøseth. Wasserstein geometry of wealth taxation: Rigid motions of the log-wealth distribution and the optimality of progressive wealth taxes. Working paper, 2026h.

Xavier Gabaix. Zipf's law for cities: An explanation. *Quarterly Journal of Economics*, 114(3): 739–767, 1999.

Xavier Gabaix, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. The dynamics of inequality. *Econometrica*, 84(6):2071–2111, 2016.

Fatih Guvenen, Gueorgui Kambourov, Burhanettin Kuruscu, Sergio Ocampo-Diaz, and Daphne Chen. Use it or lose it: Efficiency and redistributional effects of wealth taxation. *The Quarterly Journal of Economics*, 138(2):835–894, 2023. doi: 10.1093/qje/qjac047.

Fatih Guvenen, Gueorgui Kambourov, Burhanettin Kuruscu, and Sergio Ocampo. Book-value wealth taxation, capital income taxation, and innovation. Working paper, 7 June 2024. Theoretical companion to Guvenen et al. (2023). First draft: March 2021., 2024.

NOU 2022:20. Et helhetlig skattesystem. Technical report, Norges offentlige utredninger, Oslo, 2022.

Sherwin Rosen. The economics of superstars. *American Economic Review*, 71(5):845–858, 1981.

Matthew Smith, Owen Zidar, and Eric Zwick. Top wealth in America: New estimates under heterogeneous returns. *The Quarterly Journal of Economics*, 138(1):515–573, 2023. doi: 10.1093/qje/qjac033.

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