Heterogeneous Returns and Wealth Tax Neutrality: A Fokker–Planck Framework

Anders G Frøseth*

April 8, 2026

Abstract

We extend the Fokker–Planck framework of Frøseth (2026f) to populations of investors with heterogeneous, persistent return-generating ability. When the drift coefficient in the Langevin equation for log-wealth varies across investors, the proportional wealth tax remains a uniform drift shift but ceases to be neutral in the economic sense: its real incidence differs across ability types, and the stationary wealth distribution changes shape. We derive the extended Fokker–Planck equation on the joint space of log-wealth and ability, characterise the conditions under which the drift-shift symmetry breaks, and identify the consequences for asset prices and portfolio allocations. The analysis connects the neutrality results of Frøseth (2026a) and the Fokker–Planck dynamics of Frøseth (2026f) to the heterogeneous-returns literature, notably the "use-it-or-lose-it" mechanism of Guvenen et al. (2023).

1 Introduction

The distinction between a capital income tax and a wealth tax can be stated in two lines. Let a_i denote the wealth of investor i, r_i the rate of return on that wealth, and τ_k, τ_a the flat tax rates on capital income and wealth respectively. Under a capital income tax, the after-tax wealth of investor i is

$$a_i^{\text{after-tax}} = a_i + (1 - \tau_k) \cdot r_i \, a_i \, , \tag{1}$$

whereas under a wealth tax, assessed on the end-of-period market value, it is

$$a_i^{\text{after-tax}} = (1 - \tau_a) \cdot (a_i + r_i \, a_i) \, . \tag{2}$$

These two expressions, following Guvenen et al. (2023), contain the entire distinction. The income tax (1) scales only the *return* $r_i a_i$: the tax base is the flow. The wealth tax (2) scales the *entire end-of-period position* $(1 + r_i)a_i$: the tax base is the stock.[1]

When the rate of return is common across investors—$r_i = r$ for all i—the two systems are equivalent up to a rate adjustment, $\tau_a = r\tau_k/(1 + r)$. Both reduce every investor's after-tax

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

[1] The end-of-period convention is used throughout this paper series; see Frøseth (2026a), Remark 7.7, for the tax-base timing discussion and the connection to the Kruschwitz et al. (2023) arbitrage condition. In a single period the result is identical to beginning-of-period assessment, since both yield $(1 - \tau_a)(1 + r_i)\, a_i$.

wealth by the same fraction. Because the return is homogeneous, there is no structural difference between taxing the flow and taxing the stock.

But when r_i varies across investors, the equivalence breaks. The income tax burden is proportional to the individual return: a high-return investor pays more; a zero-return investor pays nothing. The wealth tax, by contrast, falls on the entire end-of-period position: even a zero-return investor pays $\tau_a \, a_i$. The consequence, over time, is that wealth migrates from low-return to high-return investors under the wealth tax—an effect absent under the income tax. This is the "use-it-or-lose-it" mechanism of Guvenen et al. (2023).

The present paper develops this observation in the continuous-time Fokker–Planck framework of Frøseth (2026a,b,f). In that framework, a proportional wealth tax enters as a uniform reduction of the drift coefficient in the Langevin equation for log-wealth, leaving the diffusion structure unchanged. This drift-shift symmetry preserves the shape of the wealth distribution and the optimal portfolio weights of all investors. A central consequence, developed in Frøseth (2026d), is the *redistribution paradox*: because the drift shift is uniform, a proportional wealth tax does not alter the shape of the stationary wealth distribution. The tax is simultaneously non-distortionary and non-redistributive through the market channel—any redistribution must come through the fiscal channel (tax revenue spent on transfers or public goods).

This paradox rests on the assumption that all investors face the same return process. Individual wealth paths differ because the Brownian increments are independent draws from a common distribution, not because investors differ in their capacity to generate returns. The resulting inequality arises from luck within a shared stochastic process. Under these conditions, the wealth tax cannot alter relative positions: it shifts the entire distribution uniformly in log-wealth.

Empirically, this assumption is well supported in deep, liquid public markets. The fund performance literature shows little evidence of persistent outperformance among active managers investing in listed equities (Fama and French, 2010). After fees, actively managed funds do not systematically beat passive benchmarks, consistent with the view that no investor has a durable edge.

However, in less efficient segments of the capital market—venture capital, private businesses, and entrepreneurial activity—there is clear evidence of persistent performance differences across managers and firms. Kaplan and Schoar (2005) find that top-quartile venture capital funds are significantly more likely to raise top-quartile successor funds. Korteweg and Sorensen (2017) confirm persistent skill net of selection effects. Most directly, Fagereng et al. (2020) document substantial persistence in individual rates of return using 20 years of Norwegian administrative panel data, finding that individual fixed effects explain a large share of return variation.

This empirical heterogeneity motivates the present paper. At the theoretical level, Benhabib et al. (2011) established that heterogeneous returns—not just earnings risk—are needed to generate realistic Pareto tails in the wealth distribution. Cao and Luo (2017) show in a general equilibrium model that persistent return heterogeneity produces a Pareto tail whose index depends on equilibrium variables, and that changes in corporate taxation and financial regulation can account for the joint evolution of rising wealth inequality and declining labour share. The continuous-time heterogeneous-agent framework of Achdou et al. (2022) provides the mathematical toolkit—Kolmogorov forward equations on the joint space of wealth and individual states—that we adopt here.

We ask: what happens to the drift-shift symmetry—and hence to the neutrality result—when investors differ in a persistent, individual-specific ability parameter that governs their expected return? The question is motivated directly by Guvenen et al. (2023), who build an overlapping-generations model with heterogeneous entrepreneurial ability and show that wealth taxation dominates capital income taxation in efficiency terms. Their "use-it-or-lose-it" mechanism—the reallocation of capital from low-ability to high-ability entrepreneurs—is absent in homogeneous-agent models and operates precisely through the differential real incidence of a uniform tax on investors with heterogeneous returns.

A central finding of the present analysis is that the common objection to wealth taxation—that it "punishes skill"—reverses the actual mechanism. Under heterogeneous returns, the wealth tax *preserves* drift differences between high- and low-ability investors; it is the income tax that compresses them. The genuine difficulty is subtler: the wealth tax cannot distinguish *why* returns are persistent. Skill and market structure interact multiplicatively—a skilled entrepreneur earns persistent excess returns only because the market is too inefficient to arbitrage the advantage away—so the tax simultaneously rewards productive ability and amplifies structural rents.

We proceed as follows. Section 2 recaps the homogeneous framework from Frøseth (2026f). Section 3 introduces the extended Fokker–Planck equation on the joint space of log-wealth and ability, derives the marginal wealth dynamics, and characterises the conditions under which the drift-shift symmetry breaks. Section 4 analyses the consequences for neutrality, equilibrium prices, and portfolio allocations. Section 5 extends the analysis to the combined flow-and-stock tax system of Frøseth (2026c), showing that the distinction between taxing the return (flow) and taxing the level (stock), which is immaterial under homogeneous returns, becomes the central question under heterogeneous returns. Section 6 discusses the empirical domain of each assumption and identifies directions for further work.

2 The homogeneous framework: recap

We briefly recall the setup from Frøseth (2026f), establishing notation for the extension that follows.

2.1 Individual dynamics

An investor's wealth $W(t)$ evolves under geometric Brownian motion:

$$\frac{\mathrm{d}W}{W} = \mu\,\mathrm{d}t + \sigma\,\mathrm{d}B_t\,, \tag{3}$$

where μ is the expected instantaneous return, $\sigma > 0$ is volatility, and B_t is a standard Brownian motion. In log-wealth $x \equiv \ln W$:

$$\mathrm{d}x = v\,\mathrm{d}t + \sigma\,\mathrm{d}B_t\,, \qquad v \equiv \mu - \frac{\sigma^2}{2}\,. \tag{4}$$

This is a Langevin equation with constant drift velocity v and additive noise of strength σ.

2.2 Population dynamics

For a population of $\mathcal{N}$ investors with common (μ, σ), the density $\pi(x,t)$ of log-wealth evolves according to the Fokker–Planck equation:

$$\frac{\partial \pi}{\partial t} = -v\frac{\partial \pi}{\partial x} + D\frac{\partial^2 \pi}{\partial x^2}\,, \qquad D \equiv \frac{\sigma^2}{2}\,. \tag{5}$$

2.3 The drift-shift symmetry

A proportional wealth tax at rate τ_w modifies only the drift:

$$\frac{\partial \pi}{\partial t} = -(v - \tau_w)\frac{\partial \pi}{\partial x} + D\frac{\partial^2 \pi}{\partial x^2}\,. \tag{6}$$

The transformation $v \mapsto v - \tau_w$ is the drift-shift symmetry. It preserves the diffusion coefficient, the shape of the propagator, and—when coupled with the asset pricing analysis of Frøseth (2026a)—the optimal portfolio weights.

The key structural requirement is that the drift shift is *uniform*: every investor experiences the same reduction τ_w in drift velocity. This uniformity is what we now relax.

3 Heterogeneous returns: the extended framework

3.1 Ability as a second state variable

We introduce a persistent, individual-specific ability parameter z that governs the investor's expected return. Specifically, replace (3) with

$$\frac{\mathrm{d}W_i}{W_i} = \mu(z_i)\,\mathrm{d}t + \sigma(z_i)\,\mathrm{d}B_t^i\,, \qquad \mu(z_i) \;\equiv\; \mathbb{E}^{\mathbb{P}}[r_i \,|\, z = z_i]\,, \tag{7}$$

where the conditional expectation is under the physical measure $\mathbb{P}$: investor i earns a different expected return because the return-generating process itself depends on ability, not because investors hold different beliefs about a common process. The volatility $\sigma(z)$ is similarly ability-dependent, and B_t^i are independent Brownian motions. We allow both drift and diffusion to depend on ability, though the case $\sigma(z) = \sigma$ (common volatility, heterogeneous drift) already captures the essential mechanism.

The ability parameter z is itself stochastic. We model it as a diffusion process:

$$\mathrm{d}z_i = -\gamma(z_i - \bar{z})\,\mathrm{d}t + \eta\,\mathrm{d}Z_t^i\,, \tag{8}$$

where $\gamma > 0$ controls the rate of mean-reversion toward the population average $\bar{z}$, $\eta > 0$ is the ability noise strength, and Z_t^i is a Brownian motion independent of B_t^i. The mean-reversion captures the empirical finding that ability persistence is real but imperfect: top-performing investors tend to revert over long horizons, particularly across generations.

Remark (Interpretation of z). The ability parameter z can represent several distinct economic mechanisms: genuine skill in identifying undervalued assets or managing enterprises; informational advantage from networks, experience, or sector knowledge; access to deal flow or co-investment opportunities unavailable to the general market; or scale advantages that larger portfolios command (lower transaction costs, better terms, access to illiquid strategies). The mathematical framework is agnostic about the source of heterogeneity; what matters is that z is persistent and affects expected returns. The empirical evidence (Fagereng et al., 2020) suggests that individual fixed effects in returns are substantial and long-lived, without cleanly identifying which mechanism dominates.

3.2 The joint Fokker–Planck equation

Define log-wealth $x_i \equiv \ln W_i$. By Itô's lemma:

$$\mathrm{d}x_i = v(z_i)\,\mathrm{d}t + \sigma(z_i)\,\mathrm{d}B_t^i\,, \qquad v(z) \equiv \mu(z) - \frac{\sigma(z)^2}{2}\,. \tag{9}$$

The state of each investor is the pair (x_i, z_i). Let $f(x, z, t)$ denote the joint density: $f(x, z, t)\,\mathrm{d}x\,\mathrm{d}z$ is the fraction of investors with log-wealth in $[x, x + \mathrm{d}x]$ and ability in $[z, z + \mathrm{d}z]$ at time t.

Since (x_i, z_i) satisfies a two-dimensional Itô diffusion with independent noise sources, the joint density evolves according to the Fokker–Planck equation:

$$\boxed{\frac{\partial f}{\partial t} = -\frac{\partial}{\partial x}[v(z)\,f] + \frac{\partial^2}{\partial x^2}[D(z)\,f] + \frac{\partial}{\partial z}[\gamma(z-\bar{z})\,f] + \frac{\eta^2}{2}\frac{\partial^2 f}{\partial z^2}\,,} \tag{10}$$

where $D(z) \equiv \sigma(z)^2/2$ is the ability-dependent diffusion coefficient.

The first two terms describe the wealth dynamics at fixed ability: an investor of ability z drifts in log-wealth at velocity $v(z)$ and diffuses with coefficient $D(z)$. The third and fourth terms describe the ability dynamics: mean-reversion toward $\bar{z}$ at rate γ and diffusion with coefficient $\eta^2/2$.

Remark (The homogeneous limit). When $\mu(z) = \mu$ and $\sigma(z) = \sigma$ for all z, the wealth dynamics decouple from ability. Integrating (10) over z recovers the homogeneous Fokker–Planck equation (5), which is the Fokker–Planck formulation of the multiplicative wealth dynamics studied by Bouchaud and Mézard (2000). The drift-shift symmetry holds in this limit because the drift is independent of the state variable z over which ability is distributed.

3.3 The taxed dynamics

Under a proportional wealth tax at rate τ_w:

$$\frac{\mathrm{d}W_i}{W_i} = (\mu(z_i) - \tau_w)\,\mathrm{d}t + \sigma(z_i)\,\mathrm{d}B_t^i\,. \tag{11}$$

In log-wealth:

$$\mathrm{d}x_i = (v(z_i) - \tau_w)\,\mathrm{d}t + \sigma(z_i)\,\mathrm{d}B_t^i\,. \tag{12}$$

The taxed joint Fokker–Planck equation is:

$$\frac{\partial f}{\partial t} = -\frac{\partial}{\partial x}[(v(z) - \tau_w)\,f] + \frac{\partial^2}{\partial x^2}[D(z)\,f] + \frac{\partial}{\partial z}[\gamma(z-\bar{z})\,f] + \frac{\eta^2}{2}\frac{\partial^2 f}{\partial z^2}\,. \tag{13}$$

The tax enters as $v(z) \mapsto v(z) - \tau_w$, exactly as in the homogeneous case: a uniform shift of the drift coefficient.

3.4 Where the symmetry breaks

The drift-shift transformation $v(z) \mapsto v(z) - \tau_w$ is still a symmetry of the Fokker–Planck *operator*: it changes only the advection term in the x-direction. But the economic content of neutrality—that the shape of the wealth distribution and the relative positions of investors are preserved—no longer follows.

To see why, consider the marginal wealth density $\pi(x,t) = \int f(x,z,t)\,\mathrm{d}z$. Integrating (13) over z:

$$\frac{\partial \pi}{\partial t} = -\frac{\partial}{\partial x}\Big[(\langle v(z)\rangle_x - \tau_w)\,\pi\Big] + \frac{\partial^2}{\partial x^2}[\langle D(z)\rangle_x\,\pi] + \underbrace{\frac{\partial}{\partial z}[\gamma(z-\bar{z})\,f] + \frac{\eta^2}{2}\frac{\partial^2 f}{\partial z^2}}_{\to 0}, \tag{14}$$

where the z-flux terms vanish upon integration over z with natural boundary conditions, and $\langle\cdot\rangle_x$ denotes the conditional expectation over ability given log-wealth x:

$$\langle v(z)\rangle_x \equiv \frac{\int v(z)\,f(x,z,t)\,\mathrm{d}z}{\pi(x,t)}\,. \tag{15}$$

The effective drift in the marginal equation is not a constant but depends on x through the conditional average $\langle v(z)\rangle_x$. If high-ability investors are concentrated at high wealth levels (a natural consequence of persistent outperformance), then $\langle v(z)\rangle_x$ is an increasing function of x. The marginal Fokker–Planck equation has *state-dependent drift*: the rich face a higher effective drift than the poor, not because the tax treats them differently, but because they are disproportionately high-ability.

Proposition 1 (Neutrality requires ability–wealth independence)**.** *The drift-shift $v(z) \mapsto v(z) - \tau_w$ preserves the marginal wealth distribution $\pi(x,t)$ if and only if the conditional average $\langle v(z)\rangle_x$ is independent of x—that is, ability and wealth are statistically independent in the joint distribution f.*

Proof. If $\langle v(z)\rangle_x = \bar{v}$ is independent of x, the marginal equation reduces to (5) with constant drift $\bar{v}$, and the shift $\bar{v} \mapsto \bar{v} - \tau_w$ is the standard drift-shift symmetry.

Conversely, if $\langle v(z) \rangle_x$ depends on x, the marginal drift is state-dependent and the shift $\langle v(z) \rangle_x \mapsto \langle v(z) \rangle_x - \tau_w$ does not preserve the shape of the stationary distribution: the ratio of drift to diffusion changes as a function of x, altering the Pareto exponent and the distributional shape. □

But ability and wealth cannot remain independent when ability is persistent and affects returns. High-ability investors accumulate faster; over time, they migrate to the upper tail. The joint distribution $f(x, z, t)$ develops a positive correlation between x and z even if the initial distribution is factored.

Proposition 2 (Endogenous correlation)**.** *Suppose $v(z)$ is strictly increasing and the initial condition is factored: $f(x, z, 0) = \pi_0(x)\, g_0(z)$. Then for any $t > 0$, the conditional mean $\langle z \rangle_x$ is a strictly increasing function of x. The wealth and ability variables are positively correlated:*

$$\mathrm{Cov}(x, z) > 0 \qquad \textit{for all } t > 0\,. \tag{16}$$

Proof sketch. At $t = 0$, all wealth levels share the same ability distribution. Over the interval $[0, \mathrm{d}t]$, investors with $z > \bar{z}$ receive drift $v(z) > \bar{v}$ and shift rightward in log-wealth relative to the mean; investors with $z < \bar{z}$ shift leftward. This differential displacement creates a positive tilt in the conditional distribution of z given x: the right tail of log-wealth becomes enriched in high-z investors, and the left tail in low-z investors. The mean-reversion in z partially offsets this tilt but cannot eliminate it as long as $v(z)$ is nonconstant and ability persistence is nonzero ($\gamma < \infty$). □

4 Consequences for neutrality, prices, and portfolios

4.1 Neutrality breaking: the mechanism

Propositions 1 and 2 together establish that the drift-shift symmetry breaks whenever returns are heterogeneous and persistent. The uniform reduction $v(z) \mapsto v(z) - \tau_w$ imposes a larger *relative* burden on low-ability investors, whose gross drift is smaller: their net drift $v(z) - \tau_w$ may turn negative while high-ability investors remain in surplus. Wealth migrates from low-z to high-z types, reshaping the stationary distribution. This is Guvenen et al.'s "use-it-or-lose-it" mechanism in Fokker–Planck language: the tax does not alter any individual's stochastic environment (it is non-distortionary in the sense of P1) but it is not neutral in the distributional sense of P3.

The redistribution paradox of Frøseth (2026d)—that a proportional wealth tax is simultaneously non-distortionary and non-redistributive—therefore holds only in the homogeneous limit. Under heterogeneous returns, the same tax actively redistributes wealth through the market channel by imposing a uniform cost on investors with heterogeneous capacities to bear it.

4.2 The effective Pareto exponent

In the homogeneous framework, the stationary Pareto exponent of the wealth distribution (when source and sink terms are included) depends on the drift-to-diffusion ratio v/D. A uniform drift shift preserves this ratio's structure and hence the Pareto exponent (Frøseth, 2026d).

With heterogeneous ability, the effective Pareto exponent of the upper tail is governed by the high-z investors who dominate the right tail:

$$\alpha_{\text{eff}} \approx 1 + \frac{v(z_{\max}) - \tau_w}{D(z_{\max})}, \tag{17}$$

where $z_{\max}$ is the ability level of the investors who populate the tail. The tax τ_w reduces the effective Pareto exponent, thinning the right tail—a redistributive effect absent in the homogeneous framework.

Remark (Dependence on the full ability distribution). Equation (17) singles out the tail-dominating ability level $z_{\max}$. In the mean-field model of Bernard et al. (2026), where agents have quenched heterogeneous growth rates drawn from a distribution with variance Σ_0^2, the stationary wealth tail exponent in the partially localised phase takes the form $\mu = 1 - \Sigma_0^2/\sigma^4$. The tail exponent thus depends on the *variance* of the growth-rate distribution, not only on its maximum. This suggests that (17) captures the leading-order effect but that the full ability distribution—in particular the dispersion of returns—enters at the next order, further thinning or thickening the tail beyond what the single dominant type predicts.

4.3 Asset prices

In the homogeneous framework, asset prices are invariant under the wealth tax because all investors are marginal for all assets, and the drift shift affects them identically. The stochastic discount factor is common.

With heterogeneous ability, the set of marginal investors for each asset may differ. If the wealth tax causes low-ability investors to shed risky assets (because they cannot sustain the tax from their lower returns), the identity of the marginal buyer changes. The remaining marginal investors—disproportionately high-ability—have different risk preferences and opportunity sets, shifting equilibrium risk premia.

Guvenen et al. (2023) find this quantitatively in their calibrated model: the equilibrium interest rate falls and the return to entrepreneurial capital rises after a tax reform from capital income taxation to wealth taxation. In the Fokker–Planck language, this corresponds to the aggregate drift $\langle v(z)\rangle_x$ adjusting endogenously as the composition of investors at each wealth level changes.

4.4 Portfolio allocations

The portfolio weight invariance of Frøseth (2026a) depends on multiplicative separability: the tax scales all asset payoffs by the same factor, so the tangent portfolio is unchanged.

When ability interacts with returns, the effective opportunity set differs across investors. A high-z investor's after-tax excess return on risky capital is $\mu(z_{\text{high}}) - \tau_w - r_f$; a low-$z$ investor's is $\mu(z_{\text{low}}) - \tau_w - r_f$. The ratio of these excess returns changes with τ_w:

$$\frac{\mu(z_{\text{high}}) - \tau_w - r_f}{\mu(z_{\text{low}}) - \tau_w - r_f} \tag{18}$$

is not invariant under shifts in τ_w (it is undefined when the denominator passes through zero). Investors optimally reallocate: low-ability investors exit risky assets; high-ability investors absorb the freed capacity. Portfolio invariance breaks not because the tax is non-proportional, but because investors face different effective opportunity sets when ability interacts with the tax.

5 Flow Taxes Versus Stock Taxes Under Heterogeneous Returns

The preceding analysis focused on the wealth tax (a stock tax) in isolation. In practice, investors face both stock and flow taxes simultaneously. Frøseth (2026c) shows that under three structural conditions—equal capital income and corporate tax rates (C1), shielding at the risk-free rate (C2), and uniform wealth tax assessment (C3)—the combined system acts through a *drift-shift-and-rescale* transformation that preserves portfolio neutrality. The flow taxes rescale excess drifts by $(1 - \tau_c)(1 - \tau_d)$; the wealth tax shifts all drifts by $-\tau_w$. Neither modification alters relative drifts between assets.

Ability heterogeneity breaks this symmetry in structurally different ways for the two tax types.

5.1 A motivating example

The following example, adapted from Guvenen et al. (2023), illustrates the asymmetry.

Consider two investors, each with wealth $a = 1000$. Investor 1 earns return $r_1 = 0\%$ (low ability); investor 2 earns $r_2 = 20\%$ (high ability). Compare a capital income tax at rate $\tau_k = 25\%$ with a wealth tax at rate $\tau_w = 2.5\%$, assessed on end-of-period wealth (consistent with equation (2)).

Table 1: After-tax wealth under capital income tax versus wealth tax, for two investors with identical wealth but different returns.

	Pre-tax	**Cap. income tax** $\tau_k = 25\%$	**Wealth tax** $\tau_w = 2.5\%$
Investor 1 ($r_1 = 0\%$)			
End-of-period wealth	1000	1000	1000
Tax paid	—	0	25
After-tax wealth	1000	1000	975
Investor 2 ($r_2 = 20\%$)			
End-of-period wealth	1200	1200	1200
Tax paid	—	50	30
After-tax wealth	1200	1150	1170
Total tax revenue	—	50	55
Wealth gap (after tax)	200	150	195

Under the capital income tax, investor 1 pays nothing (no income to tax) and investor 2 bears the entire burden. The wealth gap narrows from 200 to 150: the flow tax compresses the distribution by taxing the productive investor more heavily.

Under the wealth tax, assessed on end-of-period wealth, investor 1 pays $0.025 \times 1000 = 25$ and investor 2 pays $0.025 \times 1200 = 30$. The tax is not identical—investor 2's higher return enlarges the tax base—but the compression is marginal: the wealth gap falls from 200 to 195, compared with 150 under the income tax. The stock tax nearly preserves relative positions while the flow tax compresses them sharply. And the *composition* changes: investor 1 loses wealth in absolute terms (from 1000 to 975), while investor 2 gains (from 1000 to 1170, net of tax). Over multiple periods, this asymmetry compounds: capital migrates from low-return to high-return investors. (Under homogeneous returns, $r_1 = r_2$, the two systems are equivalent—cf. Section 1.)

5.2 The Fokker–Planck formulation

In the continuous-time framework, consider the combined tax system of Frøseth (2026c): corporate tax at rate τ_c, capital income tax at rate τ_k, dividend/gains tax at rate τ_d, and wealth tax at rate τ_w. For an investor of ability z holding a portfolio of risky assets in corporate form and a risk-free asset held personally, the after-tax dynamics in log-wealth are:

$$\mathrm{d}x_i = \Big[\underbrace{(1-\tau_c)(1-\tau_d)\big(\mu(z_i)-r_f\big)}_{\text{rescaled excess drift (flow taxes)}} + \underbrace{r_f(1-\tau_k)}_{\text{after-tax risk-free}} - \underbrace{\tau_w}_{\text{drift shift (stock tax)}} - \frac{\sigma(z_i)^2}{2}\Big]\mathrm{d}t + \sigma(z_i)\,\mathrm{d}B_t^i\,. \tag{19}$$

The flow taxes enter through the factor $(1-\tau_c)(1-\tau_d)$, which multiplies the *excess return* $\mu(z_i) - r_f$. The wealth tax enters as $-\tau_w$, a uniform shift independent of returns.

Under heterogeneous returns, these two modifications have structurally different incidence:

Flow tax incidence. The flow-tax burden on investor i is proportional to the excess return $\mu(z_i) - r_f$. A high-ability investor earns a high excess return and pays correspondingly more flow tax. A low-ability investor earns a low (or zero) excess return and pays correspondingly less. The flow tax is *ability-weighted*: its burden scales with z.

Stock tax incidence. The wealth-tax burden on investor i is τ_w per unit of log-wealth, independent of z. Both high- and low-ability investors pay the same tax per unit of wealth. But the *real* burden—the tax relative to the investor's capacity to regenerate wealth—is inversely related to ability. The stock tax is *ability-blind* in its rate but *ability-dependent* in its real incidence.

5.3 Drift decomposition under heterogeneous returns

Write the after-tax drift of investor i as:

$$v^{\text{tax}}(z_i) = \underbrace{(1-\tau_c)(1-\tau_d)}_{\equiv\lambda}\big(\mu(z_i)-r_f\big) + r_f(1-\tau_k) - \tau_w - \frac{\sigma(z_i)^2}{2}\,. \tag{20}$$

Consider two investors with abilities $z_H > z_L$. Their drift difference is:

$$v^{\text{tax}}(z_H) - v^{\text{tax}}(z_L) = \lambda\big[\mu(z_H) - \mu(z_L)\big] - \tfrac{1}{2}\big[\sigma(z_H)^2 - \sigma(z_L)^2\big]\,. \tag{21}$$

The wealth tax τ_w does not appear: it cancels in the drift *difference*. The flow taxes appear through $\lambda < 1$: they *compress* the drift difference between the two investors.

This reveals the fundamental asymmetry:

Proposition 3 (Differential incidence of flow and stock taxes)**.** *Under heterogeneous returns with $\mu(z_H) > \mu(z_L)$ and common volatility $\sigma(z) = \sigma$:*

1. *The flow taxes reduce the drift difference between high- and low-ability investors by the factor $\lambda = (1-\tau_c)(1-\tau_d)$. They compress the ability-driven wealth divergence.*

2. *The wealth tax does not affect the drift difference. It shifts all drifts uniformly, preserving the ability-driven divergence rate.*

Consequently, for a given total tax burden, a revenue-neutral shift from flow taxes toward wealth taxation widens *the drift gap between high- and low-ability investors, accelerating the reallocation of wealth from low-z to high-z types.*

Proof. Part (1) follows directly from (21): the flow-tax factor λ multiplies the return difference. Part (2) follows from the cancellation of τ_w in (21). Revenue-neutrality implies that lowering τ_c, τ_d (raising λ toward 1) requires raising τ_w. The drift difference $\lambda[\mu(z_H) - \mu(z_L)]$ increases while the uniform shift τ_w increases but does not enter the difference. □

Remark (Guvenen's three channels, reinterpreted)*.* Guvenen et al. (2023) decompose their welfare result into three channels: the "use-it-or-lose-it" reallocation, the general equilibrium price response, and the behavioral savings response. Proposition 3 provides the Fokker–Planck counterpart of the first channel: the reallocation arises because the stock tax preserves drift differences while the flow tax compresses them. A revenue-neutral switch from flow to stock taxation therefore mechanically widens the drift gap, driving the redistribution of wealth toward high-ability investors that constitutes the use-it-or-lose-it effect.

5.4 Implications for the generalised neutrality theorem

The generalised neutrality theorem of Frøseth (2026c) states that under (C1)–(C3), the combined flow-and-stock tax system preserves portfolio weights. The theorem is derived under homogeneous returns. How does it extend to heterogeneous returns?

The portfolio optimality condition for investor i depends on the after-tax excess returns *available to that investor*. Under homogeneous returns, all investors face the same excess returns (scaled by λ), and the tangent portfolio is common. Under heterogeneous returns, investor i's after-tax excess return on risky asset k relative to the risk-free asset is:

$$R_k^{\text{ex}}(z_i) = \lambda\big[\mu_k(z_i) - r_f\big] - \tau_w(\alpha_k - \alpha_0)\,, \qquad \mu_k(z_i) \;\equiv\; \mathbb{E}^{\mathbb{P}}\big[r_k \,\big|\, z = z_i\big]\,. \tag{22}$$

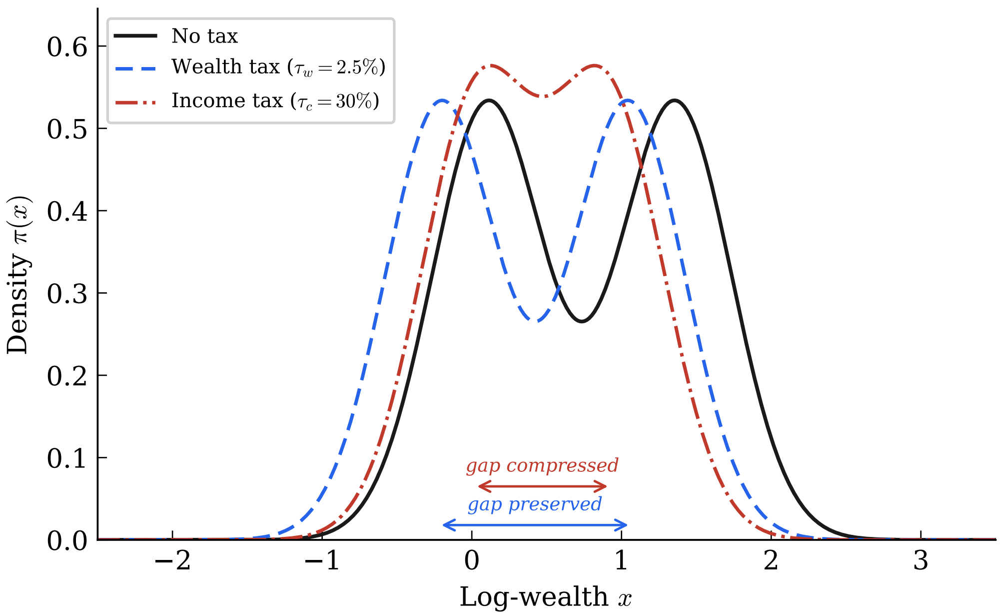

Figure 1: Stationary wealth distribution for a two-type economy (high-ability $\mu_H = 12\%$, low-ability $\mu_L = 2\%$) under three tax regimes. The wealth tax shifts both peaks leftward by the same amount, preserving the inter-type gap (Proposition 3). The income tax compresses the gap by the factor $\lambda = 1 - \tau_c$. Parameters: $\sigma = 0.15$, $\delta = 0.08$, $\tau_w = 2.5\%$, $\tau_c = 30\%$, equal population shares.

As in (7), the conditional expectation is under the physical measure $\mathbb{P}$: heterogeneity enters through conditioning on the state variable z, not through the probability measure (cf. Remark 6.4).

Under condition (C3), the wealth tax term cancels across risky assets. But the return $\mu_k(z_i)$ itself depends on investor ability. If high-ability investors earn higher returns on the *same* assets—through better deal terms, lower transaction costs, or superior selection—then different investors face different effective opportunity sets even when holding the same assets.

In the limiting case where ability affects the level of returns uniformly ($\mu_k(z) = z \cdot \mu_k$ for all k), the relative excess returns between risky assets are independent of z:

$$\frac{R_k^{\text{ex}}(z)}{R_j^{\text{ex}}(z)} = \frac{\mu_k - r_f/z}{\mu_j - r_f/z}, \tag{23}$$

which does depend on z through the effective risk-free rate r_f/z. Portfolio weights differ across ability types even under (C1)–(C3). The generalised neutrality theorem holds *within* each ability type (tax rates do not affect that type's portfolio) but the aggregate portfolio—the wealth-weighted average across types—shifts as the wealth distribution across types changes.

6 Discussion

6.1 Relationship to the existing framework

The results above do not invalidate the neutrality framework of P1–P3. Rather, they delineate its domain of validity. The drift-shift symmetry is a mathematical identity: a proportional wealth tax always enters as a uniform drift reduction. The economic content of neutrality—that the wealth distribution and relative prices are preserved—holds when this uniform shift translates into a uniform real burden across investors, which requires that ability and wealth be statistically independent (Proposition 1). When ability is persistent, Proposition 2 shows this independence cannot hold, and the symmetry breaks economically.

6.2 What the Fokker–Planck formulation adds

The joint Fokker–Planck equation (10) nests both the homogeneous and heterogeneous cases. The deviation from neutrality is controlled by a single structural feature: $\partial\langle v(z)\rangle_x/\partial x$, the conditional dependence of drift on wealth. This suggests a natural perturbation theory: when ability heterogeneity is small (low η, fast mean-reversion γ), deviations from neutrality are second-order; when heterogeneity is large, the full joint dynamics must be tracked.

6.3 Empirical implications

The framework generates a testable prediction: the degree of neutrality breaking should correlate with the share of wealth held in asset classes where return persistence is empirically documented. Within any economy, the listed-equity and government-bond segments—where informational efficiency constrains return dispersion—should exhibit approximate neutrality. The private-business, venture-capital, and entrepreneurial segments—where persistent ability differentials are well documented—should exhibit stronger distributional effects from wealth taxation. The aggregate deviation from neutrality depends on the wealth-weighted mix of these segments.

The Norwegian evidence. Fagereng et al. (2020) use Norwegian administrative panel data spanning 1993–2013—covering all taxpayers, with third-party-reported holdings across listed equities, bonds, deposits, private businesses, and real estate—to establish three facts directly relevant to the present framework.

First, individual-level return heterogeneity is large: the standard deviation of real returns on net worth is 8.6 percentage points, and the unweighted 90th–10th percentile spread exceeds 15 percentage points. Second, the heterogeneity is persistent: individual fixed effects account for approximately 26% of the variance in returns (the R^2 of the return regression rises from 0.30 without individual fixed effects to 0.50 with them), and the distribution of these fixed effects has a 90–10 percentile range of nearly 8 percentage points. Even deposit accounts—the most homogeneous asset class—exhibit a fixed-effect standard deviation of 2.6 percentage points, indicating that return heterogeneity extends beyond portfolio choice into deposit pricing and financial literacy.

Third, returns are positively correlated with wealth: moving from the 10th to the 90th percentile of net worth increases the average before-tax return by 11 percentage points, conditional on individual fixed effects and time effects (the unconditional gradient is approximately 18 percentage points).

The composition of Norwegian wealth connects these facts to the present framework. Fagereng et al. (2020) document that at the top of the wealth distribution, unlisted business equity[2] dominates: it constitutes approximately 38% of gross wealth for the top 1% and 85% for the top 0.01%.[3] Return heterogeneity exists across all asset classes—housing, listed equities, bonds—but is most extreme in unlisted businesses, where the standard deviation of returns reaches 52 percentage points and business owners exhibit a far wider distribution of individual fixed effects than non-owners. The aggregate return persistence is therefore disproportionately driven by the asset class that dominates top-percentile wealth.

This interacts with the framework in two ways. First, the large private-business share at the top implies a large effective variance of $\mu(z)$ across wealthy investors, amplifying the neutrality breaking identified in Propositions 1 and 2. Second, the wealth composition interacts with condition (C3) of Frøseth (2026c)—the requirement that the wealth tax assessment base is uniform across assets. The Norwegian *verdsettelsesrabatt* (valuation discount) assesses listed equities and unlisted shares at different fractions of market value. This breaks (C3) precisely where the return heterogeneity is concentrated: in the private-business sector. A third channel reinforces the other two: Alstadsæter et al. (2019) show that tax evasion through offshore holdings is heavily concentrated at the very top of the wealth distribution, further eroding the effective uniformity of the assessment base. The three symmetry-breaking channels—heterogeneous returns, non-uniform assessment, and differential evasion—compound rather than offset.

The US evidence. Smith et al. (2023) use administrative tax data to estimate top wealth in the United States under heterogeneous returns. They find substantial return heterogeneity within asset classes—the interest rate on fixed income at the top is approximately 3.5 times higher than the average—and document that the top 1%, 0.1%, and 0.01% wealth shares reached 33.7%, 15.7%, and 7.1% respectively by 2016. Hubmer et al. (2021) decompose the drivers of rising US wealth inequality and find that portfolio heterogeneity and the positive correlation between returns and wealth are essential to match the data—a quantitative confirmation of the endogenous correlation identified in Proposition 2.

Cross-country variation. The prediction is not that some countries exhibit neutrality and others do not, but that the degree of deviation varies with the composition of wealth at the top. The Norwegian data show that even in a Nordic economy with well-developed public equity markets, private-business wealth dominates at the very top—and it is this segment that drives the aggregate return persistence. Guvenen et al. (2023) note that despite much lower income inequality, Norwegian wealth inequality is comparable to the United States, consistent with the

[2] Fagereng et al. use the term "private equity" to denote direct ownership stakes in unlisted firms—closely held businesses, family firms, and entrepreneurial ventures—as recorded in the Norwegian shareholder registry. This is distinct from the financial-industry usage of "private equity" to denote buyout and venture capital funds.

[3] Computed from Table 1A of Fagereng et al. (2020), which reports gross-wealth composition by fractile. The 85% figure is read directly from the top-0.01% row; the 38% is the approximate population-weighted average across the three fractiles comprising the top 1%.

role of heterogeneous returns in shaping the wealth distribution. Countries where top-percentile wealth is more concentrated in listed equities and diversified portfolios should lie closer to the homogeneous benchmark. The aggregate deviation from neutrality depends on this composition and varies across economies and over time.

6.4 Sources of return persistence and the interpretation of ability

The efficiency argument for wealth taxation in Guvenen et al. (2023) rests on a specific interpretation of the ability parameter z: it reflects *productive skill* that generates genuine economic value. A high-z investor earns high returns because she selects better projects, manages firms more effectively, or deploys capital to higher-value uses. The use-it-or-lose-it mechanism is then welfare-improving: the wealth tax accelerates the reallocation of capital from less productive to more productive hands.

But the empirical finding of persistent return heterogeneity—the individual fixed effects of Fagereng et al. (2020)—is consistent with several structural mechanisms that have little to do with individual skill. The welfare implications of wealth taxation depend critically on which mechanism dominates. We identify three distinct sources of return persistence, each with different implications for the framework.

Remark (Relationship to heterogeneous beliefs). A large literature models return heterogeneity through heterogeneous beliefs: investors hold different subjective probability measures over a *common* return-generating process (Lintner, 1969; Jouini and Napp, 2006; Basak, 2005). The ability parameter z in the present framework is formally more general: each investor faces a *genuinely different* drift $\mu(z_i)$, not a different perception of a common drift. At the level of the Fokker–Planck equation, the distinction is invisible—Propositions 1–3 hold regardless of the micro-foundation. It matters, however, for three reasons: *persistence* (heterogeneous beliefs are self-correcting under learning; the structural sources identified below are persistent by construction); *welfare* (compressing belief-driven drift gaps may protect overconfident investors, while compressing ability-driven gaps destroys allocative efficiency); and *calibration* (γ should be fast under learning and slow under structural persistence—the decade-scale persistence documented by Fagereng et al. (2020) favours the latter). The framework encompasses heterogeneous beliefs as a special case but primarily targets the structural sources catalogued in Categories I–III below.

Category I: Skill in private markets. Some entrepreneurs are genuinely better at deploying capital: better judgement in project selection, superior management, deeper sector expertise. These skill differentials generate persistent return heterogeneity because private markets lack the arbitrage mechanisms—short-selling, free capital flow, public information—that enforce return homogeneity in listed equities (Section 1). Under this interpretation, z captures productive ability, and the use-it-or-lose-it mechanism is efficient. But a subtle point is already visible: skill generates persistent returns *only because* the market is inefficient. The same entrepreneur in a perfectly efficient market earns zero excess return. The persistence is a joint product of skill and market structure; we formalise this interaction below.

Category II: Structural advantage in classical private markets. Return persistence in private markets can also arise from features of market structure that are orthogonal to skill. An investor locked into a specific firm inherits that firm's trajectory—its sector exposure, its vintage, its local market position—regardless of her managerial ability. Scale economies in capital access mean that larger firms face lower borrowing costs and better terms, compounding a size advantage that originated in timing or luck rather than skill. Survivorship selection inflates apparent persistence: firms that drew unfavourable shocks exit the sample, and among survivors, returns appear more persistent than in the full population. Valuation noise in administrative data—the Norwegian tax valuations of private businesses may systematically diverge from market values—can further amplify measured return heterogeneity in precisely the asset class where it is already largest.

Under this interpretation, the individual fixed effects of Fagereng et al. (2020) partly reflect firm characteristics (sector, size, vintage, local market structure) rather than individual skill. The wealth tax's use-it-or-lose-it mechanism then redistributes capital based on structural position rather than productive ability—a reallocation that is at best neutral and at worst reinforcing of incumbency advantages.

Category III: Winner-take-all dynamics in digital markets. Rosen (1981) showed that when output can be replicated at low cost, small quality differences translate into large return differences. Digital technologies amplify this through near-zero marginal cost, global distribution, and network effects (Brynjolfsson and McAfee, 2014). The empirical consequences are documented by Autor et al. (2020) (superstar firms account for declining labour share) and De Loecker et al. (2020) (rising markups concentrated among the largest firms). In a winner-take-all market, two founders of identical ability launching identical products months apart can earn vastly different returns: the first mover captures the market; the second gets nothing. The return to the dominant firm reflects structural incumbency—network lock-in, data advantages, switching costs—not ongoing skill.

The three-way decomposition. As a first approximation, the framework's ability parameter z conflates contributions from all three sources:

$$z_i = \underbrace{z_i^{\text{skill}}}_{\text{I: productive ability}} + \underbrace{z_i^{\text{structure}}}_{\text{II: market frictions}} + \underbrace{z_i^{\text{platform}}}_{\text{III: winner-take-all rents}} . \tag{24}$$

A further refinement is needed within Category I itself. Not all persistent individual-level return heterogeneity reflects productive ability—some reflects compensation for undiversifiable risk. The fund performance literature makes this distinction precise: after controlling for systematic risk factors, the residual alpha is small in listed equities (Section 1) but significant and persistent for top-quartile venture capital. The skill component therefore admits a prior decomposition:

$$z_i^{\text{skill}} = z_i^{\alpha} + z_i^{\text{risk}} , \tag{25}$$

where z_i^{α} is genuine alpha (productive ability net of risk compensation) and z_i^{risk} is the return premium for bearing undiversifiable risk—illiquidity, concentration, vintage exposure. Define

$\varphi \in [0, 1]$ as the fraction of observed skill-component heterogeneity attributable to risk compensation. At $\varphi = 0$, all persistence reflects productive ability and the Guvenen et al. (2023) efficiency argument applies in full. At $\varphi = 1$, all persistence reflects risk bearing: the use-it-or-lose-it mechanism reallocates capital toward investors who are simply more exposed to compensated risk, with no efficiency gain. The welfare case for wealth taxation over income taxation scales with $(1 - \varphi)$, making the empirical decomposition of return persistence into alpha and risk-factor components a first-order question for tax design.

The five-component framework—z^{α}, z^{risk}, $z^{\text{structure}}$, and z^{platform} (with the structural component further split into rents and competitive frictions)—is developed in the companion literature review. For the formal results of Sections 3 to 5, the coarser three-way decomposition (24) suffices; the φ parameterisation enters through the welfare interpretation in Section 6.5 below.

The additive decomposition (24) is pedagogically useful but misleading: as noted under Category I, skill and market structure are not independent. A more faithful representation is multiplicative:

$$z_i = \theta_i \cdot h(E_j) + z_j^{\text{structure}} + z_j^{\text{platform}}, \tag{26}$$

where θ_i denotes individual ability, E_j measures the informational efficiency of the market segment j in which investor i operates, and $h(\cdot)$ is decreasing in E_j with $h(E) \to 0$ as $E \to 1$ (a fully efficient market). The first term captures the interaction: skill earns persistent excess returns only to the extent that the market fails to compete them away. The second and third terms capture rents attributable purely to market position, independent of who occupies that position.

The non-separability is empirically grounded. When high-performing professionals change institutional context, measured performance drops sharply, consistent with a large $h(E_j)$ amplifier. The Abowd et al. (1999) decomposition of wages into person effects, firm effects, and match-quality residuals confirms that the interaction term is first-order, not a perturbation.

This interaction has a natural interpretation in the Fokker–Planck framework. The mean-reversion rate γ in the Ornstein–Uhlenbeck dynamics (8) is not a property of the individual but of the market segment. In efficient markets γ is large: skill advantages are quickly arbitraged. In private markets γ is small: the same skill differentials persist because no competitive mechanism erodes them. The *observed* persistence of the skill component is θ_i/γ_j—both the numerator (ability) and the denominator (market efficiency) matter.

This non-separability has three consequences.

First, the persistence parameter γ differs systematically across the three categories, but for interacting reasons. Under Category I, γ is moderate—not because skill itself mean-reverts, but because the partial arbitrage mechanisms that do exist (imitation, entry, technology diffusion) gradually erode the advantage. Under Category II, γ is small: structural advantages—local monopolies, relationship capital, regulatory protections—face fewer erosion mechanisms. Under Category III, $\gamma^{\text{platform}} \approx 0$: once a platform achieves dominance, the dynamics are self-reinforcing. The spectral gap of the joint Fokker–Planck operator shrinks across the three categories, and the convergence time of the wealth distribution to its new steady state lengthens correspondingly.

Second, the welfare criterion changes. When the skill-times-efficiency interaction (26) dominates, the planner faces a genuine dilemma: the use-it-or-lose-it mechanism simultaneously reallocates capital toward genuine productive ability *and* toward the market segments where structural inefficiency amplifies that ability into persistent rents. When structural and platform rents dominate (the second and third terms), the planner prefers to compress drift heterogeneity—which, per Proposition 3, favours flow taxation over stock taxation. The optimal tax design therefore depends on the decomposition of observed return persistence into its components, a question that is ultimately empirical and likely sector-specific.

Third, the identification problem is genuine. An empirical decomposition exploiting the Norwegian administrative data—regressing individual return fixed effects on firm-level characteristics (sector, size, age, market concentration)—separates the individual residual from the structural component, but the individual residual itself reflects skill *as amplified by* market inefficiency. A restaurant owner in Tromsø earning persistent 15% returns may be both a talented operator and the beneficiary of a two-competitor local market; these contributions cannot be separated by conditioning on firm characteristics alone. The Abowd et al. (1999) person-firm decomposition is more informative: it isolates the portable component of returns by tracking individuals across firms. An entrepreneur whose excess returns survive a move to a different firm (controlling for firm fixed effects) is demonstrating skill that is at least partially separable from market structure. But even the person fixed effect is contaminated if the individual selects into equally inefficient markets—the portability test requires moves across markets of varying efficiency, not just across firms.

6.5 Policy implications: reinterpreting Guvenen

A common objection to wealth taxation is that it "punishes skill": by taxing the accumulated capital of successful entrepreneurs, it penalises precisely those who have demonstrated the ability to deploy capital productively. This framing reverses the actual mechanism. Proposition 3 shows that the wealth tax *preserves* skill-driven drift differences—it is the income tax that compresses them. Under heterogeneous returns, a revenue-neutral shift *from* income taxation *to* wealth taxation widens the drift gap between high- and low-ability investors, accelerating the reallocation of capital toward the skilled. This is Guvenen et al.'s central finding: the optimal wealth tax yields a welfare gain of 9.61% in consumption-equivalent terms, compared with 6.28% under optimal capital income taxation; the revenue-neutral reform (switching to a 1.13% wealth tax) delivers 7.86%. Critically, when return heterogeneity is eliminated in their robustness experiment VI ($z_{ih} = 1$), the welfare gain from switching instruments drops to 0.005%—the entire case for wealth taxation over capital income taxation rests on the heterogeneous-returns channel.

There is a genuine paradox in this result. High-ability investors accumulate more wealth and therefore pay more wealth tax in absolute terms ($\tau_w W_i$ is larger when W_i is larger)—the observation that grounds the "punishes skill" objection. But the uniform reduction $v(z) \mapsto v(z) - \tau_w$

imposes a larger *relative* burden on low-ability investors (Section 4), so wealth migrates toward the skilled over time. The absolute burden is proportional to the *stock* of wealth; the incentive structure depends on the *cross-sectional distribution* of after-tax drifts. The tax that *appears* to burden the wealthy—because it is levied on accumulated wealth—is the one that preserves the reallocation mechanism. The income tax, by contrast, may extract nothing from a zero-return investor, but it compresses the drift gap: the tax that appears gentler to skilled investors is the one that erodes the reward to skill.

The distinction between the two tax instruments connects to a deeper asymmetry in the "silent partner" metaphor. Domar and Musgrave (1944) and Stiglitz (1969) established that a proportional income tax with full loss offset acts as a silent partner on the *return margin*: the government shares proportionally in both gains and losses on $r_i - r_f$, and the investor responds by scaling up risk exposure to restore the pre-tax Sharpe ratio. The partnership operates through the return distribution. The wealth tax, as formalised in Frøseth (2026a), acts as a silent partner on the *position margin*: the government takes a proportional slice of the asset stock itself, reducing the scale of the position by $(1-\tau_w)$ each period while leaving the return distribution per unit of remaining wealth unchanged. The income-tax partnership encourages more risk-taking; the wealth-tax partnership preserves the risk composition. When returns reflect risk compensation (the z^{risk} component of (25)), the two mechanisms diverge sharply: the income tax triggers the Domar–Musgrave offset, amplifying gross risk exposure; the wealth tax leaves the portfolio structure intact. When returns reflect productive ability (z^{α}), both instruments preserve the Sharpe ratio, but only the wealth tax preserves the cross-sectional drift gap (Proposition 3). The φ parameterisation thus determines not only the welfare sign of wealth taxation but also the behavioural channel through which the two instruments differ.

The genuine difficulty is not that the wealth tax punishes skill, but that it cannot distinguish skill from other sources of return persistence. The drift gap $\mu(z_H)-\mu(z_L)$ is preserved regardless of its source ((26)):

1. Under Category I (skill amplified by market inefficiency), the wealth tax accelerates efficient reallocation. Capital moves from less productive to more productive hands. Output rises. This is the wealth tax as "capitalism's handmaiden"—strengthening the competitive allocation mechanism that is capitalism's core claim to efficiency. But the returns that the wealth tax preserves are $\theta_i \cdot h(E_j)$: the skill component is entangled with the market-inefficiency amplifier.

2. Under Categories II and III (structural advantage and winner-take-all), the wealth tax accelerates concentration. Capital moves from investors without structural advantages to those with them—from small firms to dominant platforms, from new entrants to incumbents, from competitive sectors to oligopolistic ones. The wealth tax reinforces the very market structures that competition policy seeks to counteract.

3. In the typical case—a mixture of skill and structure—the wealth tax simultaneously rewards genuine productive ability *and* amplifies structural rents. The net welfare sign depends on the ratio θ_i/γ_j: when skill is high relative to the arbitrage rate (talented en-

trepreneurs in moderately inefficient markets), the efficiency gains dominate; when structural persistence is extreme relative to skill differences (platform monopolies where the identity of the incumbent barely matters), the concentration costs dominate.

The second and third cases have implications beyond efficiency. A competitive market economy depends on a dispersed ownership structure where many independent actors compete, and where the price mechanism—not the identity of the owner—determines allocation. When the wealth tax systematically channels capital toward firms with market power, the resulting concentration translates into political influence, regulatory capture, and the erosion of the competitive discipline on which the market economy depends. The wealth tax, in this case, does not merely fail to redistribute—it actively concentrates economic power in positions that are already insulated from competitive pressure.

Proposition 3 identifies the precise policy lever. Flow taxes compress drift differences; stock taxes preserve them. If drift differences arise from productive skill, compressing them is wasteful—the Guvenen argument against income taxation. If drift differences arise from market power or structural incumbency, compressing them is *desirable*: it is the tax system doing what competition policy alone cannot, by eroding the return advantage that market structure confers. The non-separability of skill and structure means that any tax instrument necessarily affects both channels: a flow tax that compresses structural rents also compresses skill-driven returns, and a wealth tax that preserves skill-driven returns also preserves structural rents.

The implication is that the optimal mix of flow and stock taxation depends on an empirical prior: the share of observed return persistence attributable to each component of (26). In sectors where the skill-times-inefficiency interaction dominates—traditional entrepreneurship, professional services, early-stage venture capital—the wealth tax is the superior instrument despite its simultaneous preservation of market-inefficiency rents. In sectors where pure structural and platform rents dominate—platform markets, natural monopolies, industries with strong network effects—flow taxation is preferred. A uniform wealth tax applied across all sectors simultaneously gets the welfare sign right in some market segments and wrong in others.

The framework thus points toward a more nuanced instrument design than a simple uniform wealth tax: either 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 quantitative balance depends on the decomposition of return persistence in each economy—a question that the Norwegian and US administrative data are, in principle, rich enough to approximate even if a clean separation is impossible.

6.6 Open questions

Several directions remain open:

Convergence dynamics and the spectral gap. Gabaix et al. (2016) prove that, for homogeneous random growth processes with stabilising forces, the wealth distribution converges to its Pareto steady state at a rate governed by the spectral gap of the Kolmogorov forward operator. For the basic drift–diffusion process, the spectral gap is $\lambda_2 = \mu^2/(2\sigma^2)$, and convergence half-lives are typically decades to centuries for empirically plausible parameters. This slow convergence is central to P3's and P4's analysis of wealth tax dynamics.

The joint Fokker–Planck operator (10) is a two-dimensional linear operator on (x, z) space. Its spectral gap is bounded above by the minimum of two components: the wealth-dimension gap (which, following Gabaix et al.'s analysis, depends on the drift-to-diffusion ratio $v(z)/D(z)$ for the relevant ability type) and the ability-dimension gap (which depends on the mean-reversion rate γ). When ability is highly persistent (γ small), the ability dimension contributes a near-zero eigenvalue, making convergence even slower than the already-slow homogeneous case.

Moreover, Gabaix et al.'s type-dependent analysis (their Section 5) shows that when agents have different persistent growth rates μ_j, the convergence rate is stratified by type: each type j converges at its own rate $\Lambda_j(\xi) = \xi\mu_j - \xi^2\sigma_j^2/2 + \delta$, and the overall convergence is dominated by the slowest type. Translated to the continuous-ability setting of our framework: the joint distribution converges at a rate determined by the ability type with the smallest spectral gap. Low-ability investors, whose drift $v(z_L)$ is small, converge most slowly—and it is precisely these investors whose wealth dynamics are most altered by the tax.

A further subtlety is the separation of timescales: the Pareto exponent of the upper tail (a local property of high-z investors) adjusts relatively quickly to the wealth tax, while the full distribution—reflecting the slow reallocation from low-z to high-z types—adjusts at the rate of the joint spectral gap. The distributional signature of a wealth tax may therefore appear in summary statistics long before it manifests in the full distribution.

Spectral portfolio theory and aggregate allocations. The spectral portfolio theory of Frøseth (2026e) establishes that isotropic perturbations to portfolio allocation matrices—perturbations that affect all assets uniformly—preserve the spectral structure of the allocation: the tail exponent of the singular value distribution, the eigenportfolio directions, and the effective spectral rank are all invariant. A proportional wealth tax is exactly such an isotropic perturbation, and the tax neutrality corollary follows.

Under heterogeneous ability, individual-level spectral invariance still holds, but the population-level spectral structure—the wealth-weighted aggregate of all investors' allocation matrices—changes as $f(x, z, t)$ evolves under (13). High-ability investors retain their portfolio allocations while low-ability investors lose weight in the aggregate, tilting the aggregate eigenportfolios toward the factor exposures of the able. The research direction is to extend P7's spectral invariance from a single-agent result to weighted mixtures of allocation matrices whose mixing weights evolve according to the joint Fokker–Planck equation.

Mean-field feedback. In general equilibrium, asset prices depend on the wealth distribution, creating a self-consistent Fokker–Planck equation as in the McKean–Vlasov extension of Frøseth (2026d). With heterogeneous ability, the feedback structure is richer: the composition of marginal investors affects prices, which in turn affect the accumulation dynamics.

Optimal taxation. The optimal drift design framework of Frøseth (2026d) seeks the drift modification that minimises the distance to a target distribution subject to intervention costs.

With heterogeneous ability, the designer has a richer action space: taxes can be conditioned on observable proxies for ability (fund performance, firm profits) in addition to wealth. Saez and Stantcheva (2018) derive sufficient-statistics formulas for optimal capital income taxation that extend naturally to this setting. Gerritsen et al. (2025) show that a strictly positive tax on capital income is Pareto-efficient whenever rates of return are heterogeneous—whether through ability or scale effects—providing a normative counterpart to the positive analysis here. The "flat wealth tax" result of Guvenen et al. (2023)—that a uniform rate dominates ability-contingent rates—and the broader survey of wealth tax design in Scheuer and Slemrod (2021) suggest robustness properties that warrant formal investigation within the Fokker–Planck framework.

Quantitative calibration. The perturbation parameter controlling the deviation from neutrality—effectively the variance and persistence of heterogeneous returns—can in principle be estimated from the Fagereng et al. data. This would provide a quantitative measure of how far the homogeneous framework lies from the empirical truth, broken down by asset class and market segment.

References

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2022). Income and wealth distribution in macroeconomics: A continuous-time approach. *Review of Economic Studies*, 89(1):45–86.

Alstadsæter, A., Johannesen, N., and Zucman, G. (2019). Tax evasion and inequality. *American Economic Review*, 109(6):2073–2103.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. (2020). The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics*, 135(2):645–709.

Basak, S. (2005). Asset pricing with heterogeneous beliefs. *Journal of Banking & Finance*, 29(11):2849–2881.

Benhabib, J., Bisin, A., and Zhu, S. (2011). The distribution of wealth and fiscal policy in economies with finitely lived agents. *Econometrica*, 79(1):123–157.

Bernard, M., Bouchaud, J.-P., and Le Doussal, P. (2026). A mean-field theory for heterogeneous random growth with redistribution. *arXiv preprint*. arXiv:2503.23189v3.

Bouchaud, J.-P. and Mézard, M. (2000). Wealth condensation in a simple model of economy. *Physica A*, 282(3–4):536–545.

Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton.

Cao, D. and Luo, W. (2017). Persistent heterogeneous returns and top end wealth inequality. *Review of Economic Dynamics*, 26:301–326.

De Loecker, J., Eeckhout, J., and Unger, G. (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics*, 135(2):561–644.

Domar, E. D. and Musgrave, R. A. (1944). Proportional income taxation and risk-taking. *Quarterly Journal of Economics*, 58(3):388–422.

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

Fama, E. F. and French, K. R. (2010). Luck versus skill in the cross-section of mutual fund returns. *The Journal of Finance*, 65(5):1915–1947.

Frøseth, A. G. (2026a). Asset returns, portfolio choice, and proportional wealth taxation. arXiv:2603.05264 [physics.soc-ph].

Frøseth, A. G. (2026b). Extensions to the wealth tax neutrality framework. arXiv:2603.05277 [physics.soc-ph].

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

Frøseth, A. G. (2026d). From gravity to confinement: Wealth redistribution as optimal drift design in the Fokker–Planck framework. Working paper.

Frøseth, A. G. (2026e). Spectral portfolio theory: From SGD weight matrices to wealth dynamics. arXiv:2603.09006 [physics.soc-ph].

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

Gabaix, X., Lasry, J.-M., Lions, P.-L., and Moll, B. (2016). The dynamics of inequality. *Econometrica*, 84(6):2071–2111.

Gerritsen, A., Jacobs, B., Spiritus, K., and Rusu, A. V. (2025). Optimal taxation of capital income with heterogeneous rates of return. *The Economic Journal*, 135(665):180–211.

Guvenen, F., Kambourov, G., Kuruscu, B., Ocampo-Diaz, S., and Chen, D. (2023). Use it or lose it: Efficiency and redistributional effects of wealth taxation. *The Quarterly Journal of Economics*, 138(2):835–894.

Hubmer, J., Krusell, P., and Smith, A. A. (2021). Sources of US wealth inequality: Past, present, and future. In Eichenbaum, M. S. and Hurst, E., editors, *NBER Macroeconomics Annual 2020*, volume 35, pages 391–455. University of Chicago Press.

Jouini, E. and Napp, C. (2006). Heterogeneous beliefs and asset pricing in discrete time: An analysis of pessimism and doubt. *Journal of Economic Dynamics and Control*, 30(7):1233–1260.

Kaplan, S. N. and Schoar, A. (2005). Private equity performance: Returns, persistence, and capital flows. *The Journal of Finance*, 60(4):1791–1823.

Korteweg, A. and Sorensen, M. (2017). Skill and luck in private equity performance. *Journal of Financial Economics*, 124(3):535–562.

Kruschwitz, L., Löffler, A., Lorenz, D., and Uttscheid, M. (2023). Costs of capital and wealth tax: Remarks on Bjerksund and Schjelderup (2022). Discussion Paper 279, arqus.

Lintner, J. (1969). The aggregation of investor's diverse judgments and preferences in purely competitive security markets. *Journal of Financial and Quantitative Analysis*, 4(4):347–400.

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

Saez, E. and Stantcheva, S. (2018). A simpler theory of optimal capital taxation. *Journal of Public Economics*, 162:120–142.

Scheuer, F. and Slemrod, J. (2021). Taxing our wealth. *Journal of Economic Perspectives*, 35(1):207–230.

Smith, M., Zidar, O., and Zwick, E. (2023). Top wealth in America: New estimates under heterogeneous returns. *The Quarterly Journal of Economics*, 138(1):515–573.

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