

From Gravity to Confinement: Wealth Redistribution as Optimal Drift Design in the Fokker–Planck Framework

Anders G. Frøseth*

July 3, 2026

Abstract

A proportional wealth tax acts as a uniform gravitational field on the wealth distribution: it shifts the drift of the Fokker–Planck equation without altering the diffusion, preserving the Gini coefficient at all finite times. The same drift-shift symmetry that makes the tax non-distortionary (Frøseth, 2026f,h) also makes it non-redistributive through the market channel. Redistribution requires breaking this symmetry. A progressive tax (confining potential) replaces the Pareto steady state with a thinner-tailed distribution whose Gini is a closed-form function of the progressivity parameter; source-sink terms (tax-funded transfers) reshape the density directly.

We formulate optimal redistribution as a control problem for the Fokker–Planck equation, penalising intervention costs including migration, evasion, and portfolio distortion. In general equilibrium the tax design feeds back through aggregate capital and the production function, yielding a self-consistent McKean–Vlasov equation with diminishing returns to progressivity. The spectral gap of the Fokker–Planck operator determines convergence speed: progressive taxes redistribute within policy-relevant timescales, whereas proportional taxes rely on slow demographic turnover.

Keywords: Fokker–Planck equation, wealth inequality, Pareto distribution, Gini coefficient, optimal taxation, progressive wealth tax, random growth, drift-diffusion, optimal control.

JEL codes: D31, H21, H23, C61, C65.

*Independent Researcher. Email: indrefjorden@pm.me.

Contents

1	Introduction	4
2	Random growth and the Pareto steady state	7
2.1	The baseline Fokker–Planck equation	7
2.2	Demographic turnover and the Pareto tail	7
2.3	Gini coefficient and the Pareto exponent	8
3	Why proportional taxes cannot redistribute through the market channel	9
3.1	The drift-shift symmetry	9
3.2	Scale invariance of the Gini coefficient	9
3.3	The steady-state response is real but slow	10
4	From gravity to confinement: the progressive tax	11
4.1	The confining potential	11
4.2	The Ornstein–Uhlenbeck steady state	11
4.3	Gini coefficient under confinement	12
4.4	Convergence speed: the advantage of confinement	13
5	Taxonomy of Fokker–Planck modifications for redistribution	13
5.1	Class 1: Uniform drift shift (proportional tax)	14
5.2	Class 2: State-dependent drift (progressive tax)	14
5.3	Class 3: Source-sink redistribution (transfers)	14
5.4	Class 4: Absorbing boundary (wealth cap)	15
5.5	Class 5: State-dependent diffusion (volatility policy)	15
5.6	Summary	15
5.7	Leakage channels: migration and evasion	15
6	Optimal drift design	17
6.1	The control problem	17
6.2	Necessary conditions	17
6.3	The linear-quadratic case	18
6.4	The regularisation parameter as political feasibility	18
6.5	General equilibrium: the self-consistent Fokker–Planck equation	19
6.5.1	Partial versus general equilibrium: qualitative differences	19
6.5.2	The self-consistent optimal control problem	20
6.5.3	Diminishing returns to progressivity	21
7	Empirical strategy: learning the drift from data	21
7.1	The inverse problem	21

7.2	Score-based estimation	22
7.3	Neural SDE approach	22
7.4	Identifying the general equilibrium feedback	22
7.5	Natural experiment identification	22
8	Discussion	23
8.1	What the framework provides	23
8.2	The choice of target distribution	23
8.3	Heterogeneous returns and the combined tax environment	24
8.4	Limitations and extensions	25
8.5	Connection to the wealth tax programme	25
A	Gini coefficient for standard distributions	29

1 Introduction

This paper studies the redistributive consequences of wealth tax design within the Fokker–Planck framework for the wealth distribution.

The starting point is the neutrality result established in the companion papers. Frøseth (2026f) showed that a proportional wealth tax levied at market value is neutral with respect to portfolio choice: it reduces expected wealth and risk by the same proportion, leaving the risk–reward profile of every portfolio unchanged. Frøseth (2026h) reformulated this in Fokker–Planck language: the proportional tax enters as a uniform reduction of the drift coefficient, preserving the diffusion structure entirely. This *drift-shift symmetry* is the mathematical content of neutrality.

Neutrality entails a corollary for redistribution. The Gini coefficient is scale-invariant: a uniform drift shift produces a proportional scaling of all wealth levels at finite times, leaving relative inequality unchanged. The Pareto exponent of the steady-state distribution does eventually respond, but the transition takes decades to centuries (Gabaix et al., 2016), far beyond typical policy horizons. A tax that does not distort therefore does not redistribute *through the market channel*.

A proportional wealth tax does, however, redistribute through the *fiscal channel*: tax revenue funds government expenditure that benefits lower-wealth citizens disproportionately. In Fokker–Planck terms, the fiscal channel corresponds to a source-sink modification (Class 3 in Section 5), operating outside the drift-shift symmetry.

The analysis develops the *homogeneous-returns* case: all investors face a common drift v and diffusion D . This is the setting of the companion papers (Frøseth, 2026f,h). Empirical evidence shows that returns are in fact heterogeneous (Fagereng et al., 2020); Frøseth (2026d) analyses this case and shows that the redistribution paradox is specific to the homogeneous limit. The implications for the results below are discussed in Section 3.3 and §8.

The Fokker–Planck framework provides five capabilities that motivate its use throughout this paper. First, *exact symmetry results*: the drift-shift symmetry of the Fokker–Planck propagator yields both neutrality and non-redistribution as theorems, and identifies precisely which symmetry must be broken to achieve redistribution. Second, an *intrinsic taxonomy* of tax designs: every modification of the wealth process maps to a specific term in the equation (drift, diffusion, source-sink, or boundary condition), each producing a qualitatively different steady-state distribution. Third, *convergence rates* from the spectral theory of the Fokker–Planck operator: the eigenvalues yield explicit timescales for how fast inequality responds to a policy change—the calculation underlying the slow transition dynamics of Gabaix et al. (2016). Fourth, *unity of theory and estimation*: the

Fokker–Planck equation is the forward equation of the stochastic differential equation generating the wealth data, so that neural stochastic differential equation methods (Song et al., 2021) estimate exactly the drift and diffusion coefficients that appear in the optimal policy problem. Fifth, the general equilibrium extension maps to a *McKean–Vlasov* equation with established existence and uniqueness theory, rather than requiring the numerical self-consistency methods (e.g. Krusell and Smith, 1998) standard in macroeconomics.

This paper models both the market and fiscal channels and asks: for a given level of wealth tax revenue, what is the optimal tax *design*? Different tax structures correspond to different modifications of the Fokker–Planck equation, with qualitatively different effects on the steady-state wealth distribution:

- A **proportional tax** (uniform drift shift) preserves the Pareto tail. It is the analogue of a uniform gravitational field: every particle (investor) is pulled downward at the same rate, regardless of position.
- A **progressive tax** (state-dependent drift, increasing with wealth) acts as a *confining potential*—a harmonic trap in physical language. It replaces the Pareto steady state with a thinner-tailed distribution, potentially Gaussian in the Ornstein–Uhlenbeck limit. This is the Fokker–Planck realisation of Piketty’s (2014) proposal.
- **Explicit transfers** (source-sink terms in the Fokker–Planck equation) directly reshape the density by removing probability current at high wealth and injecting it at low wealth. This is the most aggressive intervention, corresponding to a universal basic income funded by a wealth tax.
- An **absorbing boundary** (wealth cap) truncates the Pareto tail entirely above a threshold, equivalent to a 100% marginal tax rate.

For each class, we derive the resulting steady-state distribution and its Gini coefficient as explicit functions of the policy parameters. We then formulate the *optimal drift design problem*: find the modification of the Fokker–Planck drift that minimises the distance between the resulting wealth distribution and a target, subject to a regularisation penalty reflecting the economic and political costs of intervention. This is a well-posed optimal control problem for a partial differential equation, drawing on techniques from both physics and applied mathematics.

The paper connects three bodies of work that have developed largely independently:

1. Piketty’s (2014) thesis that the gap between the return on capital r and economic growth g drives wealth concentration. In our framework, $r - g$ determines the drift of the Fokker–Planck equation: when $r > g$, the drift is positive and the Pareto tail fattens.

2. [Gabaix et al.’s \(2016\)](#) analysis of random multiplicative growth as the mechanism generating Pareto tails, and their finding that transition dynamics are inherently slow. This provides the baseline Fokker–Planck equation and explains why proportional taxes are ineffective on policy-relevant timescales.
3. The econophysics literature on wealth distribution ([Bouchaud and Mézard, 2000](#); [Drăgulescu and Yakovenko, 2000](#); [Chatterjee et al., 2004](#)), which applies statistical mechanics to economic dynamics but has not fully connected to the optimal taxation literature.

Empirical evidence from Norwegian wealth register data ([Fagereng et al., 2020](#)) shows that returns to wealth are increasing in wealth level, with an 18 percentage point gap between the 10th and 90th percentiles of the net worth distribution. In Fokker–Planck language, the drift $v(x)$ is an increasing function of log-wealth x —the system has a built-in anti-redistributive force that a uniform drift shift cannot overcome. [Fagereng et al. \(2025\)](#) further show that portfolio responses to equity premium changes (induced by Norwegian wealth tax reforms) are strong but slow, with adjustment taking 5–6 years—direct evidence of the relaxation timescales predicted by the Fokker–Planck framework.

A growing empirical literature exploits cross-country variation in wealth taxation to quantify behavioural responses. [Brühlhart et al. \(2022\)](#) use Swiss cantonal variation to estimate a taxable wealth elasticity of approximately 0.43, decomposed into migration, house price capitalisation, and real behavioural response. [Jakobsen et al. \(2020\)](#) exploit the 1989 Danish tax reform (a 1.2 percentage point cut in the marginal wealth tax rate) to show that wealth of moderately wealthy households increased by approximately 20% over eight years, with even larger effects for the very wealthy.¹ [Marti et al. \(2023\)](#) find that progressive wealth tax cuts in Switzerland explain 20–25% of increases in top wealth concentration. These empirical magnitudes are precisely the parameters needed to calibrate the regularisation term in our optimal control framework (Section 6).

The remainder of the paper is organised as follows. Section 2 establishes the baseline: random multiplicative growth with demographic turnover generates a Pareto-tailed steady state whose exponent is determined by the drift, diffusion, and turnover rate. Section 3 proves that a uniform drift shift (proportional tax) preserves the Gini coefficient at all finite times through the market channel, and affects the steady state only on timescales of decades to centuries. Section 4 introduces the confining potential (progressive tax) and derives the resulting steady state and Gini. Section 5 presents the full taxonomy of Fokker–Planck modifications for redistribution. Section 6 formulates the optimal drift design problem, first in partial equilibrium and then in general equilibrium where the drift

¹The 8-year treatment-on-the-treated effect on log wealth is 0.186 for couples in the exempted range and 0.312 for households unbound by the tax ceiling ([Jakobsen et al., 2020](#), Table II).

responds endogenously to the wealth distribution through the production function. Section 7 discusses the empirical strategy for estimating the current drift-diffusion structure from wealth register data, including identification of the general equilibrium feedback. Section 8 concludes with a discussion of normative choices, limitations, and connections to the broader wealth tax programme.

2 Random growth and the Pareto steady state

2.1 The baseline Fokker–Planck equation

Consider a population of individuals whose wealth $W_i(t)$ each evolves according to geometric Brownian motion:

$$\frac{dW}{W} = \mu dt + \sigma dB_t, \quad (1)$$

where μ is the expected return on capital, $\sigma > 0$ is the volatility, and B_t is a standard Brownian motion. This is the workhorse model of finance (Merton, 1969; Black and Scholes, 1973) and the starting point of the random growth literature on wealth distribution (Gabaix, 1999; Gabaix et al., 2016; Benhabib et al., 2011).

Applying Itô’s lemma to $x = \ln W$ gives the log-wealth dynamics:

$$dx = v dt + \sigma dB_t, \quad v \equiv \mu - \frac{\sigma^2}{2}. \quad (2)$$

This is a Langevin equation with constant drift v and constant diffusion coefficient $D = \sigma^2/2$. The probability density $\pi(x, t)$ of log-wealth across the population satisfies the Fokker–Planck equation:

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

Without additional mechanisms, the solution to (3) is a Gaussian that drifts to the right (when $v > 0$) and spreads diffusively:

$$\pi(x, t) = \frac{1}{\sqrt{4\pi Dt}} \exp\left(-\frac{(x - x_0 - vt)^2}{4Dt}\right). \quad (4)$$

This propagator captures individual-level dynamics well, but it does not produce the heavy-tailed wealth distributions observed empirically. A stationary distribution requires an additional mechanism.

2.2 Demographic turnover and the Pareto tail

Following Gabaix (1999) and Gabaix et al. (2016), we introduce demographic turnover: at rate $\delta > 0$, each individual is replaced by a new entrant drawn from a distribution

concentrated near the population mean. This adds a source-sink term to the Fokker–Planck equation:

$$\frac{\partial \pi}{\partial t} = -v \frac{\partial \pi}{\partial x} + D \frac{\partial^2 \pi}{\partial x^2} - \delta \pi + \delta \phi(x), \quad (5)$$

where $\phi(x)$ is the entrant distribution (normalised). The term $-\delta \pi$ removes individuals at rate δ ; the term $+\delta \phi$ inserts replacements.

Setting $\partial \pi / \partial t = 0$ and looking for solutions of the form $\pi_{\text{ss}}(x) \propto e^{-\alpha x}$ in the right tail (where $\phi(x) \approx 0$) yields the characteristic equation:

$$D\alpha^2 - v\alpha - \delta = 0. \quad (6)$$

The positive root gives the Pareto exponent:

$$\alpha = \frac{v + \sqrt{v^2 + 4D\delta}}{2D}. \quad (7)$$

Since $\pi_{\text{ss}}(x) \propto e^{-\alpha x}$ in log-wealth corresponds to $p(W) \propto W^{-(1+\alpha)}$ in wealth, this is a Pareto distribution with tail exponent α . Empirically, $\alpha \approx 1.5$ for wealth distributions in developed countries ([Gabaix, 2009](#)).

2.3 Gini coefficient and the Pareto exponent

For a pure Pareto distribution with exponent $\alpha > 1$, the Gini coefficient is:

$$\text{Gini} = \frac{1}{2\alpha - 1}. \quad (8)$$

For $\alpha = 1.5$, this gives $\text{Gini} = 0.5$. More generally, the relationship is monotonically decreasing: a larger α (thinner tail) corresponds to lower inequality.

Remark (Piketty’s $r > g$ in Fokker–Planck language). [Piketty](#)’s central thesis is that when the return on capital r exceeds economic growth g , wealth concentrates. In the Fokker–Planck framework, this translates directly: the drift $v = \mu - \sigma^2/2$ determines how fast wealth accumulates relative to the mean, and the Pareto exponent α depends on v/D and δ/D . When v is large (high $r - g$), α is large but the denominator $2D$ is also large if volatility is high. The precise balance between return, volatility, and demographic turnover determines the steady-state level of inequality. Piketty’s proposal of a progressive capital tax is, in Fokker–Planck language, a proposal to replace the uniform drift with a *confining potential*—the subject of [Section 4](#).

3 Why proportional taxes cannot redistribute through the market channel

3.1 The drift-shift symmetry

A proportional wealth tax at rate τ_w reduces the instantaneous return from μ to $\mu - \tau_w$, giving the taxed log-wealth dynamics:

$$dx = v_\tau dt + \sigma dB_t, \quad v_\tau \equiv v - \tau_w = \mu - \tau_w - \frac{\sigma^2}{2}. \quad (9)$$

The diffusion coefficient $D = \sigma^2/2$ is unchanged. The taxed Fokker–Planck equation is identical to (5) with v replaced by v_τ : a pure drift shift. This is the *drift-shift symmetry* of Frøseth (2026h).

3.2 Scale invariance of the Gini coefficient

The proportional tax scales all wealth by the same factor: $W_\tau(t) = (1 - \tau_w)^t \cdot W_0(t)$ relative to the untaxed path, where $W_0(t)$ is the wealth that would have obtained without the tax. At any finite time, the taxed wealth distribution is a rescaled version of the untaxed distribution.

Proposition 1 (Gini preservation under uniform drift shift). *Let $\pi(x, t)$ solve the Fokker–Planck equation (3) with drift v and initial condition $\pi_0(x)$, and let $\pi_\tau(x, t)$ solve the same equation with drift $v_\tau = v - \tau_w$. Then:*

$$\pi_\tau(x, t) = \pi(x + \tau_w t, t). \quad (10)$$

In wealth space, this is a multiplicative rescaling: $p_\tau(W, t) = p(W \cdot e^{\tau_w t}, t) \cdot e^{\tau_w t}$. Since the Gini coefficient is invariant under multiplicative rescaling of all wealth levels:

$$\text{Gini}[\pi_\tau(\cdot, t)] = \text{Gini}[\pi(\cdot, t)] \quad \forall t \geq 0. \quad (11)$$

Proof. Define $\tilde{\pi}(x, t) = \pi(x + \tau_w t, t)$. By direct computation:

$$\begin{aligned} \frac{\partial \tilde{\pi}}{\partial t} &= \tau_w \frac{\partial \pi}{\partial x} \Big|_{x+\tau_w t} + \frac{\partial \pi}{\partial t} \Big|_{x+\tau_w t} \\ &= \tau_w \frac{\partial \pi}{\partial x} \Big|_{x+\tau_w t} - v \frac{\partial \pi}{\partial x} \Big|_{x+\tau_w t} + D \frac{\partial^2 \pi}{\partial x^2} \Big|_{x+\tau_w t} \\ &= -(v - \tau_w) \frac{\partial \tilde{\pi}}{\partial x} + D \frac{\partial^2 \tilde{\pi}}{\partial x^2}. \end{aligned} \quad (12)$$

So $\tilde{\pi}$ solves the Fokker–Planck equation with drift $v - \tau_w = v_\tau$. Since the Gini coefficient

depends only on the shape of the Lorenz curve, which is invariant under translations in log-wealth (equivalently, multiplicative rescaling of wealth), (11) follows. $\square$

3.3 The steady-state response is real but slow

While the Gini is preserved at all finite times, the *steady-state* Pareto exponent does change:

$$\alpha_\tau = \frac{v_\tau + \sqrt{v_\tau^2 + 4D\delta}}{2D} = \frac{(v - \tau_w) + \sqrt{(v - \tau_w)^2 + 4D\delta}}{2D}. \quad (13)$$

If $\tau_w > 0$ reduces v_τ relative to v , the steady-state α_τ decreases (thicker tail, higher steady-state Gini) when the drift remains positive, and increases when the drift turns negative. But as shown by Gabaix et al. (2016), the convergence to this new steady state is governed by the spectral gap of the Fokker–Planck operator, which is of order δ (the demographic turnover rate). For realistic parameters ($\delta \approx 1/30$ per year), the half-life of the transition is:

$$t_{1/2} = \frac{\ln 2}{\Lambda}, \quad \Lambda \approx \delta + \frac{v_\tau^2}{4D} \quad (\text{for } v_\tau < 0), \quad (14)$$

where Λ denotes the spectral gap of the Fokker–Planck operator (not to be confused with the regularisation parameter λ introduced in Section 6). For a 2% tax with $\sigma = 30\%$ and $\mu = 8\%$, this gives $t_{1/2} \approx 21$ years. The distribution responds, but on timescales far beyond typical electoral cycles.

Remark (The redistribution paradox). Proposition 1 gives the market-channel result: the drift-shift symmetry that makes the proportional tax non-distortionary also makes it non-redistributive. As established in the introduction, the fiscal channel (government spending funded by tax revenue) operates as a source-sink term outside this symmetry and *is* redistributive—see Class 3 in Section 5. The paradox, then, is precise: the symmetry that makes the tax optimal for portfolio neutrality is exactly the symmetry that makes the market channel powerless. The remainder of this paper asks how to break this symmetry—through progressive design of the market channel—while preserving the fiscal channel that a proportional base already provides.

Breaking the drift-shift symmetry through progressivity (Section 4) adds a second, market-based redistribution mechanism on top of the fiscal channel.

Remark (Scope: homogeneous returns). As stated in the introduction, Proposition 1 and the redistribution paradox rely on the homogeneous-returns assumption. When persistent ability z generates an investor-specific drift $v(z)$, the uniform shift $v(z) \rightarrow v(z) - \tau_w$ imposes heterogeneous *relative* burdens: low-ability investors may face negative net drift while high-ability investors remain in positive territory. Over time, wealth reallocates from low- to high-ability types—a market-channel redistribution mechanism absent in the homogeneous case.

This “use-it-or-lose-it” effect (Guvenen et al., 2023; Frøseth, 2026d) means the redistribution paradox is specific to the homogeneous limit. Under heterogeneous returns, a proportional wealth tax *is* redistributive through the market channel—not by compressing the distribution (which requires the confining potential of Section 4), but by accelerating the reallocation of capital toward higher-ability investors. The optimal drift design framework of Section 6 remains the appropriate tool for *active* redistribution beyond this passive reallocation effect.

4 From gravity to confinement: the progressive tax

4.1 The confining potential

The proportional tax enters the Fokker–Planck equation as a uniform drift shift $v \rightarrow v - \tau_w$ —a constant external force, analogous to a uniform gravitational field. Gravity shifts the distribution without confining it.

A *progressive* wealth tax, where the effective rate increases with wealth, introduces a *state-dependent* drift:

$$v(x) = v_0 - \kappa(x - \bar{x}), \quad (15)$$

where $v_0 = \mu - \sigma^2/2 - \tau_w^{(0)}$ is the drift at the reference log-wealth $\bar{x}$, $\tau_w^{(0)}$ is the base tax rate, and $\kappa > 0$ is the *progressivity parameter*. The term $-\kappa(x - \bar{x})$ is a *restoring force*: it pushes high-wealth individuals ($x > \bar{x}$) leftward more strongly than low-wealth individuals.

In physics, this is a *harmonic potential*: the force increases linearly with displacement from the reference point. The resulting dynamics are those of an Ornstein–Uhlenbeck (OU) process:

$$dx = [v_0 - \kappa(x - \bar{x})] dt + \sigma dB_t. \quad (16)$$

4.2 The Ornstein–Uhlenbeck steady state

The Fokker–Planck equation for (16) is:

$$\frac{\partial \pi}{\partial t} = -\frac{\partial}{\partial x} \{ [v_0 - \kappa(x - \bar{x})] \pi \} + D \frac{\partial^2 \pi}{\partial x^2}. \quad (17)$$

Setting $\partial \pi / \partial t = 0$ and solving gives the steady-state distribution:

$$\pi_{\text{ss}}(x) = \frac{1}{\sqrt{2\pi \cdot D/\kappa}} \exp\left(-\frac{(x - \bar{x} - v_0/\kappa)^2}{2D/\kappa}\right). \quad (18)$$

This is a *Gaussian* in log-wealth, centred at $x^* = \bar{x} + v_0/\kappa$ with variance $\sigma_{\text{ss}}^2 = D/\kappa$.

The transformation is dramatic: the Pareto tail (power-law decay) is replaced by a Gaussian tail (exponential-quadratic decay). In wealth space, the steady state becomes *log-normal*—a distribution with much thinner tails than the Pareto.

4.3 Gini coefficient under confinement

For a log-normal distribution with log-variance $\sigma_{\text{ss}}^2 = D/\kappa$, the Gini coefficient is:

$$\text{Gini}_{\text{OU}} = 2\Phi\left(\frac{\sigma_{\text{ss}}}{\sqrt{2}}\right) - 1 = 2\Phi\left(\sqrt{\frac{D}{2\kappa}}\right) - 1, \quad (19)$$

where Φ is the standard normal CDF.

This is a monotonically decreasing function of the progressivity parameter κ : more progressivity confines the distribution more tightly, reducing the Gini.

Proposition 2 (Gini as a function of progressivity). *For the Ornstein–Uhlenbeck wealth dynamics (16) with diffusion coefficient $D = \sigma^2/2$ and progressivity $\kappa > 0$:*

- (i) $\text{Gini}_{\text{OU}} \rightarrow 0$ as $\kappa \rightarrow \infty$ (complete equality in the limit of infinite progressivity).
- (ii) $\text{Gini}_{\text{OU}} \rightarrow 1$ as $\kappa \rightarrow 0^+$ (the distribution spreads without bound, approaching maximal inequality).
- (iii) For a target Gini $G^* \in (0, 1)$, the required progressivity is:

$$\kappa^* = \frac{D}{2\left[\Phi^{-1}\left(\frac{1+G^*}{2}\right)\right]^2}. \quad (20)$$

Proof. Part (i): As $\kappa \rightarrow \infty$, $\sigma_{\text{ss}}^2 = D/\kappa \rightarrow 0$, so the log-normal concentrates at a point and $\text{Gini} \rightarrow 0$. Part (ii): As $\kappa \rightarrow 0^+$, $\sigma_{\text{ss}} \rightarrow \infty$ and $\text{Gini} \rightarrow 1$. Part (iii): Solve $2\Phi(\sqrt{D/(2\kappa)}) - 1 = G^*$ for κ . $\square$

Remark (Numerical illustration). For typical parameters ($\sigma = 0.30$, so $D = 0.045$), achieving the current empirical Gini of ≈ 0.80 requires $\kappa \approx 0.014$; reducing the Gini to 0.50 (the level of a pure exponential distribution, à la Drăgulescu and Yakovenko 2000) requires $\kappa \approx 0.050$; and reaching the Scandinavian income Gini of ≈ 0.27 would require $\kappa \approx 0.19$. These values correspond to marginal tax rate increases of approximately 1.4%, 5.0%, and 19% per unit of log-wealth, respectively.




Figure 1: Convergence to target wealth distribution under different tax structures. The proportional tax (blue, solid) converges at the demographic turnover rate δ ; the progressive tax (red, dashed) converges at the progressivity rate κ ; a more strongly progressive tax (black, dotted) converges faster still. The shaded region marks a typical electoral cycle (4–8 years). The progressive tax achieves meaningful redistribution within policy-relevant timescales; the proportional tax does not.

4.4 Convergence speed: the advantage of confinement

The OU process has a well-defined spectral gap equal to κ : all deviations from the steady state decay at rate $e^{-\kappa t}$. The half-life of convergence is:

$$t_{1/2}^{\text{OU}} = \frac{\ln 2}{\kappa}. \quad (21)$$

For $\kappa = 0.05$, this is $t_{1/2} \approx 14$ years—comparable to the demographic turnover timescale, but now driven by the *tax itself* rather than by births and deaths. The confining potential actively compresses the distribution, rather than passively waiting for demographic replacement.

This is a qualitative difference from the proportional tax. Figure 1 compares the convergence trajectories: the proportional tax relies on demographic turnover (slow), while the progressive tax creates its own convergence mechanism (fast).

5 Taxonomy of Fokker–Planck modifications for redistribution

We now present the full classification of policy tools by their Fokker–Planck modification, steady-state distribution, and effect on inequality.

5.1 Class 1: Uniform drift shift (proportional tax)

FP modification: $v \rightarrow v - \tau_w$ (constant shift).

Steady state: Pareto with modified exponent α_τ given by (13).

Gini: Preserved at finite times; slow convergence to $1/(2\alpha_\tau - 1)$.

Physics analogue: Uniform gravitational field.

5.2 Class 2: State-dependent drift (progressive tax)

FP modification: $v \rightarrow v_0 - \kappa(x - \bar{x})$ (linear restoring force).

Steady state: Log-normal (Gaussian in log-wealth) with variance D/κ .

Gini: $2\Phi(\sqrt{D/(2\kappa)}) - 1$, continuously adjustable via κ .

Physics analogue: Harmonic potential / Ornstein–Uhlenbeck confinement.

Convergence rate: κ (actively driven by the tax).

5.3 Class 3: Source-sink redistribution (transfers)

FP modification: Add $-\gamma(x)\pi + S(x)$ where $\gamma(x)$ is a wealth-dependent extraction rate and $S(x)$ is a transfer injection.

Steady state: Depends on the specific forms of γ and S . For a flat UBI funded by progressive extraction, the steady state acquires a lower bound (minimum wealth) and a thinner upper tail.

Gini: Can be reduced to any target level by choosing γ and S appropriately.

Physics analogue: Particle removal and injection; coupled reservoirs.

Note: This class captures the *fiscal channel* of wealth taxation. Any wealth tax—including a proportional one—generates revenue that funds government expenditure. If that expenditure benefits citizens at lower wealth levels disproportionately (through public services, welfare transfers, or infrastructure), the net effect is a source-sink modification of the wealth distribution. In this sense, even the neutral proportional tax of Class 1 is redistributive when combined with its fiscal counterpart: the drift shift preserves the market-channel Gini, but the tax-and-transfer flow reduces the *total* Gini. The revenue from a 2% wealth tax on billionaires, as proposed by [Saez and Zucman \(2019\)](#) and [Zucman \(2024\)](#), would fund substantial source terms at low x .

Table 1: Taxonomy of Fokker–Planck modifications for redistribution. Each row represents a class of policy intervention, characterised by the modification it introduces into the Fokker–Planck equation and its effect on the steady-state wealth distribution.

Policy class	FP modification	Steady state	Gini effect
Proportional tax	$v \rightarrow v - \tau_w$	Pareto (shifted α)	Preserved (finite t)
Progressive tax	$v(x) = v_0 - \kappa x$	Log-normal	Controlled by κ
Transfers (UBI)	Source–sink terms	Bounded below	Strongly reduced
Wealth cap	Absorbing boundary	Truncated Pareto	Hard upper bound
Volatility policy	$D \rightarrow D(x)$	Thinner tail	Indirectly reduced

5.4 Class 4: Absorbing boundary (wealth cap)

FP modification: $\pi(x^*, t) = 0$ for $x > x^*$, with probability current at x^* recycled to lower wealth.

Steady state: Truncated distribution with hard upper bound.

Gini: Bounded above by $\text{Gini}_{\text{Pareto}}(x^*)$, the Gini of a Pareto distribution truncated at x^* .

Physics analogue: Absorbing boundary with particle recycling.

5.5 Class 5: State-dependent diffusion (volatility policy)

FP modification: $D \rightarrow D(x)$ with $D(x)$ decreasing at high wealth.

Steady state: Thinner tail than Pareto (the tail exponent increases where D decreases).

Gini: Reduced, but the policy instrument is indirect (reducing risk for the wealthy, e.g. through portfolio insurance or diversification mandates).

Physics analogue: Spatially varying temperature; cooler regions have less fluctuation.

5.6 Summary

5.7 Leakage channels: migration and evasion

The five classes above describe what the policymaker *intends* to do to the Fokker–Planck equation. In practice, two channels cause the *realised* modification to differ from the intended one.

Migration (permeable boundary). If agents can relocate to avoid the tax, the confining potential $v(x) = v_0 - \kappa(x - \bar{x})$ operates only up to some threshold x_m beyond

which agents exit the jurisdiction. In Fokker–Planck language, this replaces the reflecting boundary at infinity with a *partially absorbing boundary*:

$$J(x_m, t) = \gamma \pi(x_m, t), \quad (22)$$

where J is the probability current and $\gamma > 0$ is a migration rate that increases with the tax burden at x_m . Jakobsen et al. (2024) document significant migration responses to wealth taxation; Pichet (2007) estimates that the French ISF caused approximately €200B in capital flight since the tax’s creation in 1988. The consequence is a truncated steady state: the Pareto tail is cut off not by policy design (Class 4) but by agent exit, and the effective Gini reduction is smaller than the closed-boundary prediction.

Tax evasion and avoidance (attenuated drift). If agents can hide or reclassify wealth, the intended progressivity κ is not fully realised. Let $\epsilon(x)$ denote the evasion rate—the fraction of wealth at position x that escapes the tax base. The effective drift modification becomes:

$$\delta v_{\text{eff}}(x) = -[1 - \epsilon(x)] \kappa (x - \bar{x}), \quad (23)$$

so the effective spring constant is $\kappa_{\text{eff}}(x) = [1 - \epsilon(x)] \kappa < \kappa$. Alstadsæter et al. (2019) estimate $\epsilon \approx 0.25$ for the top 0.01%; Bjørneby et al. (2023) document avoidance through corporate restructuring that creates a wedge between intended and realised drift. Since $\epsilon(x)$ is typically increasing in x (richer agents have more avoidance opportunities), the effective confining force weakens precisely where it is most needed.

Debt and leverage (amplified assessment distortion). When the wealth tax base is net wealth (assets minus liabilities) and assets are assessed below market value while debt is deducted at face value, the tax base diverges from market net wealth. An investor with assets of market value A , assessment ratio $\beta < 1$, and debt D has taxable net wealth $W_{\text{tax}} = \beta A - D$ versus market net wealth $W_{\text{market}} = A - D$. As shown in Frøseth (2026h), the ratio $W_{\text{tax}}/W_{\text{market}}$ decreases with leverage and can become negative. This creates an incentive to borrow against underassessed assets (notably real estate), which affects the Fokker–Planck dynamics in two ways: (i) the effective drift depends on the leverage ratio, not only on wealth level, coupling the debt decision to the wealth dynamics; and (ii) leveraged positions amplify the effective volatility of net wealth, making the diffusion coefficient D state-dependent through an endogenous mechanism. Unlike evasion, which attenuates the intended drift, the leverage channel can *amplify* the deviation between the observed taxable distribution and the true market-value distribution.

Remark (Leakage in the taxonomy). Migration modifies the *boundary condition* of the Fokker–Planck equation; evasion modifies the *drift coefficient*; debt leverage modifies both the *drift* (through the leverage ratio) and the *diffusion* (through amplified volatil-

ity). All three reduce the effective redistribution below the designed level, but through mathematically distinct channels. The diffusion model approach of Section 7 captures all three implicitly: the learned drift $\hat{v}(x)$ reflects the *realised* drift including evasion and leverage effects, and the learned boundary behaviour reflects migration.

6 Optimal drift design

6.1 The control problem

We now formulate the redistribution problem as an optimal control problem for the Fokker–Planck equation. The state is the wealth density $\pi(x, t)$; the control is the drift modification $\delta v(x)$; the objective is to steer the distribution toward a target $\pi^*(x)$ with prescribed inequality level.

Definition 1 (Optimal drift design problem). Given a baseline Fokker–Planck equation with drift v and diffusion D , a target distribution π^* , a time horizon T , and a regularisation parameter $\lambda > 0$, find the drift modification $\delta v(x)$ that minimises:

$$\mathcal{J}[\delta v] = \frac{1}{2} \int_{-\infty}^{\infty} |\pi(x, T) - \pi^*(x)|^2 dx + \frac{\lambda}{2} \int_{-\infty}^{\infty} |\delta v(x)|^2 dx, \quad (24)$$

subject to the controlled Fokker–Planck equation:

$$\frac{\partial \pi}{\partial t} = -\frac{\partial}{\partial x} \{ [v + \delta v(x)] \pi \} + D \frac{\partial^2 \pi}{\partial x^2} - \delta \pi + \delta \phi, \quad (25)$$

with initial condition $\pi(x, 0) = \pi_0(x)$.

The first term in (24) measures how close the resulting distribution is to the target; the second term penalises the magnitude of the intervention. The regularisation parameter λ captures the trade-off between redistributive effectiveness and the economic and political costs of intervention.

6.2 Necessary conditions

Applying Pontryagin’s maximum principle, the optimal drift modification satisfies:

$$\delta v^*(x) = -\frac{1}{\lambda} \pi(x, T) \cdot p(x, T), \quad (26)$$

where $p(x, t)$ is the adjoint variable solving the backward equation:

$$-\frac{\partial p}{\partial t} = [v + \delta v(x)] \frac{\partial p}{\partial x} + D \frac{\partial^2 p}{\partial x^2} - \delta p, \quad (27)$$

with terminal condition $p(x, T) = \pi(x, T) - \pi^*(x)$.

This is a coupled forward-backward PDE system: the forward equation propagates the density π ; the backward equation propagates the adjoint p ; and the optimal control δv^* depends on both. Standard iterative methods (gradient descent on δv , or forward-backward splitting) converge to the solution.

6.3 The linear-quadratic case

When the target distribution is log-normal (corresponding to a Gaussian in log-wealth), and we restrict the control to the linear class $\delta v(x) = -\kappa(x - \bar{x}) - c$ for constants κ, c , the problem reduces to choosing κ and c optimally. The solution follows from Proposition 2: for a given target Gini G^* , the optimal progressivity is κ^* given by (20), and c adjusts the mean. This is the linear-quadratic regulator (LQR) for the Fokker–Planck equation.

6.4 The regularisation parameter as political feasibility

The parameter λ in (24) deserves interpretation. A large λ penalises intervention heavily, yielding a near-proportional (near-neutral) tax. A small λ permits aggressive redistribution. In practice, λ encodes:

- **Distortion costs:** The welfare loss from portfolio distortion and behavioural responses (*cf.* the distortion channels of Frøseth, 2026b). Empirical estimates from Switzerland (Brühlhart et al., 2022) and Denmark (Jakobsen et al., 2020) provide calibration targets.
- **Leakage costs:** The migration and evasion channels formalised in Section 5.7 reduce the effective progressivity below the designed level. The permeable boundary (22) truncates the tail through agent exit; the attenuated drift (23) weakens the confining force. These channels do not simply increase costs—they *modify the control problem itself*: the effective control is $\kappa_{\text{eff}} = [1 - \epsilon] \kappa$, so the optimality condition (26) must be evaluated at the realised drift, not the legislated one.
- **Administrative costs:** Progressive taxes require valuation of all assets at market value, which may be infeasible for illiquid assets.
- **Political constraints:** The political economy of taxation limits how progressive the tax can be.

The optimal control framework makes this trade-off explicit and quantifiable, rather than leaving it to qualitative argument. In particular, the leakage channels imply that the policymaker must over-design κ relative to the closed-system optimum to compensate for the attenuation by $\epsilon(x)$ and the probability current lost through migration at x_m .

6.5 General equilibrium: the self-consistent Fokker–Planck equation

The analysis so far treats the pre-tax drift v as an exogenous parameter. This is the partial equilibrium assumption: the return on capital r (and hence $v = r - \sigma^2/2$) does not respond to the tax. But in general equilibrium, r depends on the aggregate capital stock K , which is itself determined by the wealth distribution π . The pre-tax return and economic growth g are jointly determined by technology, preferences, and policy—a point emphasised by [Jones \(2015\)](#) in the context of Piketty’s $r > g$ thesis.

To close the model, we introduce a production function $Y = F(K, L)$ with diminishing returns to capital. The marginal product of capital gives the pre-tax return:

$$r = F_K(K, L), \quad K[\pi] = \int_0^\infty W p(W) dW = \int_{-\infty}^\infty e^x \pi(x) dx, \quad (28)$$

where $K[\pi]$ is a functional of the wealth density. The growth rate g may also depend on K (through capital deepening), so both r and g respond endogenously to changes in π .

The drift becomes a functional of the distribution:

$$v[\pi] = r(K[\pi]) - \frac{\sigma^2}{2}. \quad (29)$$

The self-consistent Fokker–Planck equation is then:

$$\frac{\partial \pi}{\partial t} = -\frac{\partial}{\partial x} \left\{ [v[\pi] + \delta v(x)] \pi \right\} + D \frac{\partial^2 \pi}{\partial x^2} - \delta \pi + \delta \phi, \quad (30)$$

where the drift $v[\pi]$ is itself determined by π through (28)–(29). This is a *nonlinear* Fokker–Planck equation: the drift depends on the solution.

Remark (Mean-field structure). Equation (30) has the structure of a *McKean–Vlasov* or mean-field equation, familiar from statistical physics: each particle (investor) experiences a drift that depends on the aggregate state of the population. The self-consistency condition $v = v[\pi]$ is analogous to the self-consistent field in Hartree–Fock theory or the mean-field approximation in spin systems. [Bernard et al. \(2026\)](#) solve this class of model explicitly for the case of heterogeneous growth rates with mean-field redistribution, identifying a phase transition between localised (condensed) and delocalised (spread) wealth distributions at a critical redistribution intensity.

6.5.1 Partial versus general equilibrium: qualitative differences

The GE feedback introduces a countervailing force. In partial equilibrium, a confining potential with progressivity κ monotonically reduces the Gini: more κ always means

less inequality. In general equilibrium, increasing κ compresses the wealth distribution, reducing aggregate capital K . By diminishing returns, $r = F_K(K, L)$ rises, increasing the baseline drift v . This partially offsets the confinement:

$$\kappa \uparrow \Rightarrow K \downarrow \Rightarrow r \uparrow \Rightarrow v \uparrow \Rightarrow \text{partial offset of confinement} . \quad (31)$$

Remark (Neutrality and economic growth). The feedback chain (31) rests on the assumption that wealth taxation reduces productive capital accumulation. But Frøseth (2026f) showed that a proportional wealth tax at market value is *neutral* with respect to portfolio choice: asset prices, risk-taking, and portfolio allocations are unchanged. If the progressive tax is constructed as a neutral proportional base plus a confining surcharge, the base component does not distort capital allocation. Only the progressive surcharge $-\kappa(x - \bar{x})$ introduces distortions. This means the channel from taxation to lower g operates only through the progressive component, not through the full tax rate. The neutral base reduces the *level* of wealth but preserves the *composition* of risk-taking that drives innovation and growth—weakening the standard objection that wealth taxes discourage entrepreneurship. In the self-consistent framework, this implies that the GE feedback (31) may be substantially smaller than a naïve analysis (which attributes the full tax rate to capital destruction) would suggest.

The steady state is determined by the fixed point where the confining force and the GE feedback balance. For a Cobb–Douglas production function $Y = K^a L^{1-a}$ with capital share a , the fixed-point condition is:

$$v^* = a \left(K[\pi_{\kappa, v^*}^*] \right)^{a-1} L^{1-a} - \frac{\sigma^2}{2} , \quad (32)$$

where π_{κ, v^*}^* is the steady state of the FP equation with drift $v^* - \kappa(x - \bar{x})$. This is a scalar fixed-point equation in v^* (since K is a deterministic functional of the Gaussian steady state), which can be solved by simple iteration.

6.5.2 The self-consistent optimal control problem

The optimal drift design problem (Definition 1) generalises to the GE setting by replacing the exogenous drift with the self-consistent drift:

$$\mathcal{J}_{\text{GE}}[\delta v] = \frac{1}{2} \int |\pi(x, T) - \pi^*(x)|^2 dx + \frac{\lambda}{2} \int |\delta v(x)|^2 dx , \quad (33)$$

subject to (30). The key difference is that the forward equation is now nonlinear (through $v[\pi]$), so the adjoint equation acquires an additional term from the functional derivative

$\delta v / \delta \pi$. The optimality condition becomes:

$$\delta v^*(x) = -\frac{1}{\lambda} \left[\pi(x, T) \cdot p(x, T) + \int \frac{\partial v}{\partial K} \cdot e^{x'} \pi(x', T) p(x', T) dx' \right], \quad (34)$$

where the second term captures the GE feedback: the optimal policy accounts for its own effect on the aggregate return.

6.5.3 Diminishing returns to progressivity

The GE feedback implies *diminishing returns to progressivity*. Define the GE Gini reduction per unit of progressivity:

$$\eta(\kappa) \equiv -\frac{d\text{Gini}_{\text{GE}}(\kappa)}{d\kappa}. \quad (35)$$

In partial equilibrium, η is always positive (more progressivity always helps). In general equilibrium, $\eta(\kappa)$ is positive for small κ but decreasing, and may approach zero for large κ if the rise in r fully offsets the confinement. Whether η actually reaches zero—whether there is a *maximum achievable redistribution*—depends on the elasticity of the production function. For Cobb–Douglas, the fixed point always exists and $\eta > 0$ for all finite κ , but the marginal gain shrinks.

This has a direct policy implication: the GE feedback provides a natural economic limit on redistribution that is *separate* from the political constraints encoded in λ . Even with $\lambda = 0$ (no political cost), the production function constrains how much inequality can be reduced through wealth taxation alone.

7 Empirical strategy: learning the drift from data

The framework above delivers testable predictions and closed-form policy formulas, but treats the drift v and diffusion D as known. In practice, they must be estimated from data. This section outlines the estimation strategy; empirical results using wealth register data will follow in a companion paper.

7.1 The inverse problem

Norwegian wealth register data (Fagereng et al., 2020, 2025) provides individual-level wealth panels spanning two decades, covering the full population.

The estimation problem is: given observations $\{W_i(t_k)\}$ for individuals $i = 1, \dots, N$ at times $t_0, t_1, \dots, t_K$, recover $v(x)$ and $D(x)$ nonparametrically.

7.2 Score-based estimation

A score-based diffusion model learns the score function $s(x, t) = \partial \log \pi(x, t) / \partial x$ from data. At steady state, the score function is related to the FP coefficients by:

$$s_{ss}(x) = \frac{v(x)}{D(x)} - \frac{D'(x)}{D(x)}. \quad (36)$$

If D is constant (or independently estimated), the score function directly recovers the drift:

$$v(x) = D \cdot s_{ss}(x). \quad (37)$$

7.3 Neural SDE approach

An alternative is to parameterise $v(x)$ and $D(x)$ as neural networks and train them by maximum likelihood on the panel data, using differentiable SDE solvers (Kidger, 2021). This directly learns the drift and diffusion functions without going through the score.

7.4 Identifying the general equilibrium feedback

The self-consistent FP equation (Section 6.5) predicts that the drift v responds endogenously to changes in the wealth distribution. A diffusion model trained on panel data spanning a tax reform captures this feedback *implicitly*: the learned drift before the reform reflects the old equilibrium; the drift after reflects the new one, including whatever adjustment in r occurred through the aggregate capital channel.

Concretely, let $\hat{v}_{pre}(x)$ and $\hat{v}_{post}(x)$ denote the drift functions estimated from data before and after a reform that changes the tax rate by $\Delta\tau$. In partial equilibrium, the drift change would be a uniform shift: $\hat{v}_{post} - \hat{v}_{pre} = -\Delta\tau$ for all x . In general equilibrium, the aggregate capital response adds a positive offset (since K falls and r rises):

$$\hat{v}_{post}(x) - \hat{v}_{pre}(x) = -\Delta\tau + \Delta r (K[\pi_{post}]), \quad (38)$$

where $\Delta r > 0$ is the GE feedback. This is testable: if the empirically estimated drift shift is *smaller* than the mechanical tax change $\Delta\tau$, the GE feedback is present. The magnitude of Δr identifies the elasticity of the production function.

7.5 Natural experiment identification

The Norwegian wealth tax reforms of 1992 and 1998 (Fagereng et al., 2025) provide exogenous shocks to the equity premium that can be used for identification. A diffusion model trained on data spanning these reforms should detect changes in the drift coefficient,

and the *form* of the change (uniform shift versus state-dependent) tests the neutrality hypothesis of Frøseth (2026f).

Additional natural experiments are available from other countries. The 1989 Danish wealth tax reform (Jakobsen et al., 2020) provides a clean before–after comparison for wealth accumulation dynamics. Swiss cantonal variation (Brülhart et al., 2022; Marti et al., 2023; Burgherr, 2021) offers cross-sectional identification of how different tax rates affect the stationary distribution. Norwegian corporate restructuring responses to valuation discounts (Bjørneby et al., 2023) identify avoidance channels corresponding to the attenuated drift of (23): the wedge between intended and realised drift modification is directly observable. Finally, Alstadsæter et al. (2019) show that tax evasion at the top of the distribution is substantial (approximately 25% of tax liability for the top 0.01%), implying that the observed $\pi(x, t)$ from tax registers underestimates true concentration and that the evasion rate $\epsilon(x)$ of Section 5.7 is empirically large.

8 Discussion

8.1 What the framework provides

The Fokker–Planck approach to redistribution offers four advantages over existing frameworks. First, it provides a *unified language* for comparing tax designs: proportional, progressive, transfers, and caps all correspond to specific modifications of the same equation. Second, it makes the *speed–magnitude trade-off* explicit: the spectral gap of the FP operator determines how fast the distribution converges to its new steady state. Third, it connects the *micro-level* dynamics (individual wealth accumulation) to the *macro-level* outcome (the wealth distribution) through a single PDE. Fourth, it models *both redistribution channels*—market (drift modification) and fiscal (source-sink terms)—within the same equation, and the general equilibrium extension (Section 6.5) captures their interaction through the production function: the tax design affects aggregate capital, which feeds back into both market returns and the revenue that funds the fiscal channel.

8.2 The choice of target distribution

The optimal drift design problem (Definition 1) requires specifying a target distribution π^* . The mathematical framework is agnostic about this choice—it can accommodate any target—but the choice itself is ultimately one of normative political philosophy, not physics.

The framework does, however, make the consequences of different normative positions mathematically precise. A utilitarian planner maximising total expected utility would choose π^* to maximise $\int u(e^x)\pi^*(x) dx$ for some concave utility function u , subject to the

production constraint. A Rawlsian planner, reasoning from behind the *veil of ignorance* (Rawls, 1971)—uncertain of their own position in the wealth distribution—would instead maximise the welfare of the worst-off individual. In the Fokker–Planck framework, the veil of ignorance corresponds precisely to the initial condition: an agent who does not yet know which realisation of the stochastic process they will experience. The Rawlsian maximin criterion then selects the drift modification $\delta v(x)$ that maximises the minimum wealth level in the steady-state distribution, subject to the GE constraint that the production function must still be satisfied.

These are different objective functionals in the same optimal control problem. The utilitarian criterion typically yields a less aggressive confining potential than the Rawlsian criterion, because it values the surplus of the wealthy (through the concave u) whereas the maximin criterion does not. The framework makes this trade-off quantitative: for each normative choice, it delivers the corresponding optimal progressivity κ^* , convergence timescale $t_{1/2}$, and GE feedback magnitude Δr .

This paper does not advocate for a particular normative position. Its contribution is to provide the mathematical machinery that translates any such position into a concrete policy design, and to make the costs and timescales of that design explicit.

8.3 Heterogeneous returns and the combined tax environment

Two companion papers relax the homogeneous-returns assumption (Section 3.3) in ways that qualify—and in some cases strengthen—the results above.

Frøseth (2026d) introduces persistent heterogeneity in investor ability, generating drift $v(z)$ that depends on an individual parameter z . Because a proportional tax preserves drift *differences*—the quantity $v(z_H) - v(z_L)$ is unchanged by a uniform shift—it accelerates the reallocation of capital toward higher-ability investors. This “use-it-or-lose-it” mechanism (Guvenen et al., 2023) means the proportional tax *is* redistributive through the market channel under heterogeneous returns, even though the redistribution paradox holds in the homogeneous limit.

The confining potential framework of Section 4 remains the appropriate tool for *active* redistribution of the wealth distribution. But under heterogeneous returns, the baseline against which the confining potential operates is itself redistributive: the proportional tax component already generates reallocation, and the progressive surcharge adds compression on top. The optimal progressivity κ^* derived in Section 6 should therefore be interpreted as the additional confinement needed beyond the passive reallocation effect.

Frøseth (2026c) extends the neutrality framework to combined flow-and-stock taxation, showing that corporate and dividend taxes enter the Fokker–Planck equation as a rescaling

of excess drift differences by the factor $(1 - \tau_c)(1 - \tau_d)$, while the wealth tax enters as a uniform shift. The two effects are additively separable under natural institutional conditions. For the taxonomy of Section 5, this means that in a combined tax system, Class 2 modifications (progressive wealth tax) interact with the flow-tax environment: the effective drift gap between asset classes depends on both the progressivity parameter κ and the flow-tax rates τ_c, τ_d . In particular, a revenue-neutral shift from flow taxes to wealth taxes *widens* ability-driven drift gaps (Frøseth, 2026d), a consideration that enters the optimal design problem when the tax mix is itself a policy variable. That combined-instrument design problem is developed in companion work: by Frøseth (2026g) under homogeneous returns, and by Frøseth (2026e) as the heterogeneous-returns synthesis.

8.4 Limitations and extensions

Several important extensions remain for future work. The OU model (linear confinement) is a simplification; realistic progressive tax schedules have nonlinear rate structures that produce non-Gaussian steady states. Frøseth (2026g) develops two such extensions—smooth log-progressive and bracket schedules—within the combined flow-stock framework; fully general schedule classes remain open. The leakage channels formalised in Section 5.7—migration as a permeable boundary (22) and evasion as attenuated drift (23)—are treated here as exogenous. A full analysis would endogenise γ and $\epsilon(x)$ as functions of the tax design itself, coupling the leakage rates to the optimal control problem. The experience of the French ISF, where Pichet (2007) estimated substantial capital flight, and of the Swiss cantons, where Brülhart et al. (2022) documented significant behavioural responses, confirms that leakage is empirically important.

The empirical estimation of $v(x)$ and $D(x)$ from Norwegian data is the natural next step, and will determine whether the GBM baseline is an adequate starting point or whether state-dependent coefficients are essential. Zucman (2024)’s proposal for a coordinated global minimum wealth tax on billionaires would, by closing migration channels, effectively set $\gamma \rightarrow 0$ in (22) and permit stronger confinement than national policies alone.

8.5 Connection to the wealth tax programme

This paper is part of a series studying wealth taxation through the Fokker–Planck framework. Frøseth (2026f) established that a proportional market-value wealth tax is neutral with respect to portfolio choice. Frøseth (2026b) identified the channels through which neutrality breaks in practice. Frøseth (2026h) reformulated both results in Fokker–Planck language, revealing the drift-shift symmetry as the mathematical content of neutrality. Frøseth (2026c) extended the neutrality framework to combined flow-and-stock taxation, showing additive separability of flow-tax and stock-tax distortions under natural insti-

tutional conditions. Frøseth (2026d) introduced heterogeneous investor ability, showing that the proportional tax is simultaneously non-distortionary *and* redistributive when returns vary across investors—qualifying the redistribution paradox established here to the homogeneous case.

The present paper develops the optimal drift design framework: it characterises when and how the drift-shift symmetry must be broken to achieve market-channel redistribution, formulates the design problem as optimal control of the Fokker–Planck equation, and models the general equilibrium interaction between the market and fiscal channels.

Four further companions build directly on the design framework developed here. Frøseth (2026a) characterises minimum-distortion tax designs on the proportional class in closed form, contrasting an information-theoretic with a transport-geometric optimality criterion. Frøseth (2026g) lifts the progressive drift-design problem to the combined flow–stock environment under Gini and top-share criteria, including bracket schedules. Frøseth (2026e) extends that synthesis to heterogeneous returns. Frøseth (2026i) recasts the neutral class and progressive optimality in 2-Wasserstein geometry—the transport-geometric counterpart of the drift-design view taken here.

A companion empirical paper will estimate the drift and diffusion functions from Norwegian wealth register data using neural stochastic differential equation methods, testing the theoretical predictions developed here—including the neutrality hypothesis, the general equilibrium feedback, and the leakage channels—against observed wealth dynamics.

Acknowledgements

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

References

- Annette Alstadsæter, Niels Johannesen, and Gabriel Zucman. Tax evasion and inequality. *American Economic Review*, 109(6):2073–2103, 2019. doi: 10.1257/aer.20172043.
- Jess Benhabib, Alberto Bisin, and Shenghao Zhu. The distribution of wealth and fiscal policy in economies with finitely lived agents. *Econometrica*, 79(1):123–157, 2011.
- 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.

- Marie Bjørneby, Simen Markussen, and Knut Røed. An imperfect wealth tax and employment in closely held firms. *Economica*, 90(358):557–583, 2023. doi: 10.1111/ecca.12456.
- Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
- Jean-Philippe Bouchaud and Marc Mézard. Wealth condensation in a simple model of economy. *Physica A*, 282(3–4):536–545, 2000.
- Marius Brülhart, Jonathan Gruber, Matthias Krapf, and Kurt Schmidheiny. Behavioral responses to wealth taxes: Evidence from Switzerland. *American Economic Journal: Economic Policy*, 14(4):111–150, 2022. doi: 10.1257/pol.20200258.
- David Burgherr. The costs of net wealth taxation: Evidence from a tax reform in Zurich. *ETH Zurich, Public Finance Discussion Paper*, 2021.
- Arnab Chatterjee, Bikas K. Chakrabarti, and S. S. Manna. Pareto law in a kinetic model of market with random saving propensity. *Physica A*, 335:155–163, 2004.
- Adrian Drăgulescu and Victor M. Yakovenko. Statistical mechanics of money. *European Physical Journal B*, 17(4):723–729, 2000.
- Andreas Fagereng, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. Heterogeneity and persistence in returns to wealth. *Econometrica*, 88(1):115–170, 2020.
- Andreas Fagereng, Luigi Guiso, and Marius A. K. Ring. How much and how fast do investors respond to equity premium changes? evidence from wealth taxation. Working paper, 2025.
- Anders G. Frøseth. Minimum-distortion wealth taxation, i: Information-theoretic versus transport-geometric optimality on the proportional class. Working paper, 2026a.
- Anders G. Frøseth. Extensions to the wealth tax neutrality framework. 2026b. arXiv:[2603.05277](#) [physics.soc-ph].
- Anders G. Frøseth. Flow taxes, stock taxes, and portfolio choice: A generalised neutrality result. 2026c. arXiv:[2603.15974](#) [physics.soc-ph].
- Anders G. Frøseth. Heterogeneous returns and wealth tax neutrality: A Fokker–Planck framework. 2026d. arXiv:[2603.16006](#) [physics.soc-ph].
- Anders G. Frøseth. Progressive wealth taxation under heterogeneous returns: A synthesis. Working paper, 2026e.
- Anders G. Frøseth. Asset returns, portfolio choice, and proportional wealth taxation. 2026f. 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, 2026g.
- Anders G. Frøseth. Wealth taxation as a drift modification: A Fokker–Planck approach to tax neutrality. 2026h. arXiv:[2603.05283](https://arxiv.org/abs/2603.05283) [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, 2026i.
- Xavier Gabaix. Zipf’s law for cities: An explanation. *Quarterly Journal of Economics*, 114(3):739–767, 1999.
- Xavier Gabaix. Power laws in economics and finance. *Annual Review of Economics*, 1: 255–294, 2009. doi: 10.1146/annurev.economics.050708.142940.
- 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, and Daphne Chen. Use it or lose it: Efficiency and redistributive effects of wealth taxation. *Quarterly Journal of Economics*, 138(2):835–894, 2023. doi: 10.1093/qje/qjac047.
- Katrine Jakobsen, Kristian Jakobsen, Henrik Kleven, and Gabriel Zucman. Wealth taxation and wealth accumulation: Theory and evidence from Denmark. *The Quarterly Journal of Economics*, 135(1):329–388, 2020. doi: 10.1093/qje/qjz032.
- Katrine Jakobsen, Henrik Kleven, Jonas Kolsrud, Camille Landais, and Mathilde Muñoz. Taxing top wealth: Migration responses and their aggregate economic implications. NBER Working Paper No. 32153, February 2024, revised January 2026; forthcoming, *American Economic Review*, 2024.
- Charles I. Jones. Pareto and piketty: The macroeconomics of top income and wealth inequality. *Journal of Economic Perspectives*, 29(1):29–46, 2015.
- Patrick Kidger. On neural differential equations, 2021. D.Phil. thesis, University of Oxford.
- Per Krusell and Anthony A. Smith, Jr. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896, 1998.
- Samira Marti, Isabel Z. Martínez, and Florian Scheuer. Does a progressive wealth tax reduce top wealth inequality? evidence from Switzerland. *Oxford Review of Economic Policy*, 39(3):513–529, 2023.

- Robert C. Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *Review of Economics and Statistics*, 51(3):247–257, 1969. doi: 10.2307/1926560.
- Eric Pichet. The economic consequences of the French wealth tax. *La Revue de Droit Fiscal*, 14:1–25, 2007.
- Thomas Piketty. *Capital in the Twenty-First Century*. Harvard University Press, Cambridge, MA, 2014.
- John Rawls. *A Theory of Justice*. Harvard University Press, Cambridge, MA, 1971.
- Emmanuel Saez and Gabriel Zucman. Progressive wealth taxation. *Brookings Papers on Economic Activity*, 2019(Fall):437–533, 2019. doi: 10.1353/eca.2019.0017.
- Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations (ICLR)*, 2021.
- Gabriel Zucman. A blueprint for a coordinated minimum effective taxation standard for ultra-high-net-worth individuals. Report commissioned by the Brazilian G20 Presidency, 2024.

A Gini coefficient for standard distributions

For reference, we collect the Gini coefficient formulas for the distributions appearing in the main text.

Pareto distribution with tail exponent $\alpha > 1$:

$$\text{Gini}_{\text{Pareto}} = \frac{1}{2\alpha - 1} . \quad (39)$$

Log-normal distribution with log-variance σ_x^2 :

$$\text{Gini}_{\text{LN}} = 2\Phi\left(\frac{\sigma_x}{\sqrt{2}}\right) - 1 = \text{erf}\left(\frac{\sigma_x}{2}\right) . \quad (40)$$

Exponential distribution (Boltzmann–Gibbs, following [Drăgulescu and Yakovenko, 2000](#)):

$$\text{Gini}_{\text{exp}} = \frac{1}{2} . \quad (41)$$

Truncated Pareto distribution with exponent $\alpha > 1$ and upper bound W_{max} , to

leading order in $u = (W_{\min}/W_{\max})^{\alpha-1}$:

$$\text{Gini}_{\text{trunc}} = \text{Gini}_{\text{Pareto}}[1 - u] + o(u) , \quad (42)$$

so truncation reduces the Gini below its untruncated value by the fraction u at leading order.