Wealth Taxation as a Drift Modification: A Fokker–Planck Approach to Tax Neutrality

Anders G Frøseth*

Draft — 2 March 2026

Abstract

We reformulate the neutral wealth tax framework of Frøseth (2026b) in the language of stochastic dynamics and statistical physics. Individual wealth under geometric Brownian motion satisfies a Langevin equation with multiplicative noise; the probability density of wealth across a population then evolves according to a Fokker–Planck equation. A proportional wealth tax at market value enters as a uniform reduction of the drift coefficient, preserving the diffusion structure and all relative probability currents. This drift-shift symmetry is the physical content of tax neutrality. Each channel through which neutrality breaks down in practice—book-value assessment, liquidity frictions, forced dividend extraction, migration, and market impact—corresponds to a specific violation of this symmetry: a state-dependent, asset-dependent, or flow-dependent modification of the Fokker–Planck equation. The framework clarifies when wealth taxation is a benign rescaling of the dynamics and when it introduces genuinely new physics.

Contents

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

1 Introduction

The wealth tax debate, both in the academic literature and in public policy, typically proceeds in one of two registers. Economists analyse the tax through equilibrium models—CAPM, Modigliani–Miller, consumption-based asset pricing—where the object of interest is the effect on prices, returns, and portfolio weights. Separately, a large econophysics literature studies wealth distributions using tools from statistical mechanics—Boltzmann–Gibbs ensembles, Fokker–Planck equations, agent-based models—but rarely connects to the asset-pricing framework that governs how taxes interact with financial markets.

This paper bridges the two. We show that the neutrality results derived in Frøseth (2026b) and the distortion channels analysed in Frøseth (2026a) have a natural and precise formulation in the language of stochastic dynamics. The mapping is not metaphorical: the same geometric Brownian motion that underlies the finance results *is* a Langevin equation, and the evolution of the wealth distribution across investors *is* governed by a Fokker–Planck equation. No new physics is introduced; we simply read the existing mathematics in a different dialect.

The value of this exercise is threefold. First, the Fokker–Planck formulation makes the symmetry content of neutrality transparent: the wealth tax is neutral if and only if it enters as a uniform, state-independent shift of the drift coefficient. Second, each violation of neutrality maps to a specific, classifiable modification of the Fokker–Planck equation—offering a taxonomy of distortions that is both exhaustive and physically intuitive. Third, the framework naturally accommodates questions about wealth *distributions* and their evolution—relaxation times, steady states, and the interplay between individual dynamics and aggregate outcomes—that lie outside the scope of representative-agent asset pricing.

We proceed step by step. Section 2 introduces the core concepts of modern asset pricing—the stochastic discount factor, no-arbitrage, and the price of risk—in a language accessible to physicists. Section 3 establishes the core mapping: individual wealth dynamics under geometric Brownian motion as a Langevin equation with multiplicative noise. Section 4 derives the corresponding Fokker–Planck equation for the wealth distribution. Section 5 introduces the proportional wealth tax and shows that it enters as a pure drift shift. Section 6 formalises the neutrality result as a symmetry of the Fokker–Planck equation and shows that this symmetry is robust to non-Gaussian returns and stochastic volatility. Section 7 maps each distortion channel to a specific symmetry-breaking mechanism. Section 8 discusses steady-state distributions when income and consumption are included as source and sink terms. Section 9 discusses extensions and open questions.

2 Asset pricing for physicists: Cochrane's framework

This section introduces the core ideas of modern asset pricing in a language accessible to readers trained in physics or physical chemistry. The presentation follows Cochrane (2005), who showed that the entire theory reduces to a single equation and a single object—the *stochastic discount factor*. Readers familiar with finance may skip to Section 3.

2.1 What is an asset?

An *asset* is anything that produces an uncertain future payoff. A share of stock pays dividends and can be sold; a bond pays coupons and returns its face value; a house provides rent (or imputed rent) and can be resold. The central question of asset pricing is: *given the uncertain future payoff, what is the correct price today?*

In physics terms, an asset is a random variable x_{t+1} representing the total payoff (cash flow plus resale value) received at time $t + 1$. The price p_t is the value assigned to this random variable today. The *gross return* is

$$R_{t+1} \equiv \frac{x_{t+1}}{p_t}\,, \tag{1}$$

so that $R_{t+1} = 1.05$ means a 5% gain. The return is a dimensionless ratio—the payoff per unit of price—and is the financial analogue of a growth factor.

2.2 The fundamental equation of asset pricing

Cochrane's central result is that, under very weak assumptions (no arbitrage, discussed below), there exists a random variable $m_{t,t+1}$—the *stochastic discount factor* (SDF)—such that the price of *every* asset satisfies

$$\boxed{p_t = \mathbb{E}_t\big[m_{t,t+1}\, x_{t+1}\big]\,.} \tag{2}$$

Here $\mathbb{E}_t[\cdot]$ denotes the conditional expectation given all information available at time t. The equation says: the price today equals the expected value of tomorrow's payoff, weighted by m.

Remark (Analogy with statistical mechanics). Equation (2) has the same mathematical structure as a *partition-function average*. In statistical mechanics, the expectation value of an observable A in a canonical ensemble is

$$\langle A \rangle = \sum_s A(s)\, \frac{e^{-\beta E(s)}}{Z}\,, \tag{3}$$

where $\beta = 1/k_B T$ is the inverse temperature and $Z = \sum_s e^{-\beta E(s)}$ is the partition function. The Boltzmann weight $e^{-\beta E(s)}/Z$ plays exactly the role of m: it assigns a positive weight to each state, and the observable's "price" (its equilibrium expectation) is the weighted average over states.

The analogy is not perfect—m is stochastic and need not take the exponential Boltzmann form—but the structural parallel is deep: both frameworks assign state-dependent weights to compute expectations. In finance, "bad states" (recessions, crises) receive high weight because investors value payoffs more when they are poor; in physics, low-energy states receive high weight because they are thermodynamically favoured.

2.3 The stochastic discount factor

The SDF $m_{t,t+1}$ encodes how the market values payoffs in different states of the world. Its key properties:

Positivity. Under no-arbitrage (see below), $m > 0$ in all states. This is the financial analogue of the requirement that Boltzmann weights are positive.

High in bad states. When the economy is in a bad state (low consumption, high unemployment), m is large: the market assigns high value to payoffs received precisely when they are most needed. This is the origin of risk premia.

Low in good states. Conversely, payoffs received when everyone is already wealthy are worth less per unit.

In consumption-based models, m has the explicit form

$$m_{t,t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma}, \tag{4}$$

where $\beta < 1$ is a time-preference (patience) parameter, C_t is aggregate consumption, and $\gamma > 0$ is relative risk aversion. When consumption falls ($C_{t+1} < C_t$), the ratio $(C_{t+1}/C_t)^{-\gamma}$ is large, so m is large: payoffs in recessions are highly valued. But Cochrane's key insight is that (2) holds *regardless* of whether we assume (4) or any other specific model. The SDF exists as long as there is no arbitrage.

2.4 No arbitrage: the financial conservation law

No arbitrage means that there is no trading strategy that produces a positive expected payoff with zero cost and zero risk. In physics language: *there is no perpetual motion machine in financial markets.*

Proposition 1 (Fundamental theorem of asset pricing)**.** *The following are equivalent:*

1. *There are no arbitrage opportunities.*
2. *There exists a strictly positive stochastic discount factor $m > 0$ such that $p = \mathbb{E}[m\, x]$ for all traded assets.*

This is the most important theorem in finance. It says that the absence of free lunches is *equivalent* to the existence of a positive weighting function m. The parallel to physics is:

Physics	**Finance**
No perpetual motion (2nd law)	No arbitrage
$\Rightarrow$ Entropy is a state function	$\Rightarrow$ SDF exists ($m > 0$)
$\Rightarrow$ Boltzmann weights well-defined	$\Rightarrow$ Prices are expectations under m

Just as the second law of thermodynamics is not derived from microscopic dynamics but constrains all possible dynamics, no-arbitrage is not derived from individual behaviour but constrains all possible prices.

2.5 Risk, return, and the price of risk

Dividing (2) by p_t and using $R_{t+1} = x_{t+1}/p_t$ gives the *return form*:

$$1 = \mathbb{E}_t[m_{t,t+1}\, R_{t+1}] \,. \tag{5}$$

Expanding using the covariance identity $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y] + \mathrm{Cov}(X,Y)$:

$$\mathbb{E}_t[R_{t+1}] = \frac{1}{\mathbb{E}_t[m_{t,t+1}]} - \frac{\mathrm{Cov}_t(m_{t,t+1}, R_{t+1})}{\mathbb{E}_t[m_{t,t+1}]} \,. \tag{6}$$

The first term is the *risk-free rate*: the return on an asset whose payoff does not fluctuate. Since a risk-free asset has $\mathrm{Cov}(m, R_f) = 0$:

$$R_f = \frac{1}{\mathbb{E}[m]} \,. \tag{7}$$

The second term is the *risk premium*: assets that pay off in bad states (when m is high) have $\mathrm{Cov}(m,R) > 0$, hence *lower* expected returns—they provide insurance. Assets that pay off in good states (when m is low) have $\mathrm{Cov}(m,R) < 0$, hence *higher* expected returns—they carry risk.

Remark (Risk is not variance)*.* A central insight that physicists often find surprising: in finance, risk is *not* measured by the variance of an asset's return. An asset can be highly volatile yet command no risk premium if its fluctuations are uncorrelated with the SDF (i.e., uncorrelated with the aggregate state of the economy). Risk is entirely about *covariance with the pricing kernel*—not individual variance. This is analogous to the fact that, in statistical mechanics, the contribution of a mode to the free energy depends on its coupling to the thermal bath, not on its amplitude alone.

2.6 The risk-free rate as temperature

The risk-free rate $R_f = 1/\mathbb{E}[m]$ sets the baseline price of time: how much the market discounts future payoffs purely for waiting, with no risk. It plays a role analogous to temperature in statistical mechanics:

- **Low** R_f (low interest rates) $\Leftrightarrow$ $\mathbb{E}[m]$ is high $\Leftrightarrow$ future payoffs are valuable $\Leftrightarrow$ asset prices are high. This is a "hot" market in physics language: high prices, low discount rates, willingness to pay for future outcomes.
- **High** R_f (high interest rates) $\Leftrightarrow$ $\mathbb{E}[m]$ is low $\Leftrightarrow$ future payoffs are less valuable $\Leftrightarrow$ asset prices are low. A "cold" market: investors demand compensation for waiting.

The analogy is suggestive but should be used with care: unlike thermodynamic temperature, R_f is not a state variable of the system—it is an endogenous price determined by aggregate patience and growth expectations.

2.7 Specific models as specifications of m

Different finance models are simply different choices for the form of m:

Model	**SDF specification**	**Physics analogue**
CAPM	$m = a - b\,R_{\text{market}}$	Single-mode coupling
Consumption CAPM	$m = \beta(C_{t+1}/C_t)^{-\gamma}$	Boltzmann weight
Fama–French 3-factor	$m = a - b_1 f_1 - b_2 f_2 - b_3 f_3$	Three-mode coupling
Black–Scholes	$m = e^{-r_f \Delta t}$ (risk-neutral)	Constant weight

The power of Cochrane's approach is that (2) holds *before* choosing a specific model. This is analogous to deriving results from thermodynamic identities (which hold for any system) rather than from a specific Hamiltonian.

2.8 Cochrane's framework and the wealth tax

With this machinery, we can state the wealth tax neutrality result in Cochrane's language. A proportional wealth tax at rate τ_w reduces every asset's payoff by the same factor: $x_{t+1} \to (1 - \tau_w)\,x_{t+1}$. Since m is unchanged (the tax does not alter the state of the economy or the marginal utility of consumption), the price becomes

$$p_t^{\text{after tax}} = \mathbb{E}_t\big[m_{t,t+1}\,(1 - \tau_w)\,x_{t+1}\big] = (1 - \tau_w)\,\mathbb{E}_t\big[m_{t,t+1}\,x_{t+1}\big] = (1 - \tau_w)\,p_t\,. \qquad (8)$$

Every price falls by the same factor $(1 - \tau_w)$. The *return* $R_{t+1} = x_{t+1}/p_t$ is unchanged because both numerator and denominator scale by $(1 - \tau_w)$. Expected returns, risk premia, Sharpe ratios, and optimal portfolio weights are all invariant.

This is the pricing side of the neutrality result. The remainder of the paper shows that the same result has a natural expression in Fokker–Planck language: the tax is a uniform drift shift that preserves the structure of the stochastic dynamics.

Remark (When does the SDF change?). The neutrality result assumes that m is unaffected by the tax. This holds in partial equilibrium (the tax does not change aggregate consumption or the marginal utility structure) and in the general equilibrium of a homogeneous economy where all investors are taxed identically. If the tax is non-uniform—applying only to some investors or some assets—then m may change, and neutrality breaks down. This connects to the distortion channels of Section 7.

2.9 Correspondence of concepts

Table 1 collects the key conceptual correspondences between the finance and statistical physics frameworks developed in this paper. The table is intended as a reference; each entry is introduced and justified in the sections that follow.

Table 1: Correspondence between finance and statistical physics concepts. The left column lists standard finance terminology (following Cochrane 2005); the right column gives the statistical physics equivalent used in this paper. Entries above the mid-rule are conceptual; entries below are mathematical objects.

Finance	**Statistical physics**
Entities and structure	
Investor	Particle (trajectory in wealth space)
Population of investors	Ensemble
Asset (stock, bond, ...)	Degree of freedom
Portfolio of N assets	System of N coupled degrees of freedom
Market	Open system (with sources and sinks)
Prices and constraints	
No arbitrage	No perpetual motion (2nd law)
Stochastic discount factor m	Boltzmann weight / measure on states
Asset price $p = \mathbb{E}[m\,x]$	Partition-function average $\langle A \rangle$
Risk-free rate $R_f = 1/\mathbb{E}[m]$	Inverse temperature $\beta = 1/k_B T$ *(suggestive, not exact)*
Dynamics	
Expected return μ	Drift velocity $v = \mu - \sigma^2/2$
Volatility σ	Noise strength (diffusion coeff.$^{1/2}$)
GBM $\mathrm{d}W/W = \mu\,\mathrm{d}t + \sigma\,\mathrm{d}B_t$	Langevin equation (multiplicative noise)
Log-return process	Langevin equation (additive noise)
Diffusion coefficient	$D = \sigma^2/2$ (Einstein relation)
Covariance matrix $\mathbf{\Sigma}$	Coupling matrix between modes
Eigenvectors of $\mathbf{\Sigma}$	Normal modes of the system
Distributions and equilibrium	
Wealth distribution $p(W,t)$	Probability density (FP equation)
Probability current J	Net flux in wealth space
Efficient market (risk-neutral drift)	Detailed balance under pricing measure
Wealth inequality (Pareto tail)	Power-law tail of steady-state distribution
Gini coefficient, Pareto exponent	Moments / exponents of $\pi_{\mathrm{ss}}(x)$
Taxation	
Proportional wealth tax τ_w	Uniform drift shift: $v \to v - \tau_w$
Tax neutrality	Drift-shift symmetry (uniform external field)
Book-value assessment	Anisotropic field (asset-dependent coupling)
Liquidity friction	State-dependent friction (v, D depend on W)
Forced dividend extraction	Coupling to slow variable (firm capital K)
Migration (tax-induced emigration)	Absorbing boundary / sink term
Market impact of forced sales	Mean-field interaction (nonlinear drift)

3 Individual wealth as a Langevin equation

3.1 Geometric Brownian motion in financial language

Consider an investor whose entire wealth $W(t)$ is invested in a single risky asset. Under the standard assumption of geometric Brownian motion (GBM), the wealth evolves as

$$\frac{\mathrm{d}W}{W} = \mu\,\mathrm{d}t + \sigma\,\mathrm{d}B_t\,, \tag{9}$$

where μ is the expected instantaneous return, $\sigma > 0$ is the volatility, and B_t is a standard Brownian motion. This is the starting point of the Black–Scholes–Merton framework, the Capital Asset Pricing Model, and the neutrality analysis in Frøseth (2026b).

The key property of (9) is that the noise is *multiplicative*: the magnitude of the random shock $\sigma W\,\mathrm{d}B_t$ is proportional to the current wealth level. Rich investors experience larger absolute fluctuations than poor investors, even though the percentage fluctuation σ is the same.

3.2 Log-wealth and the Langevin equation

Define $x(t) \equiv \ln W(t)$. By Itô's lemma,

$$\mathrm{d}x = \underbrace{\left(\mu - \frac{\sigma^2}{2}\right)}_{\equiv\, v}\mathrm{d}t \;+\; \sigma\,\mathrm{d}B_t\,. \tag{10}$$

This is a *Langevin equation with additive noise* in the log-wealth coordinate x. The drift velocity $v = \mu - \sigma^2/2$ is constant, and the noise strength σ is independent of the state. In the physics of Brownian motion, (10) describes a particle drifting at constant velocity v in a viscous medium subject to thermal fluctuations of strength σ (Zwanzig, 2001, Ch. 2).

Remark (No restoring force). There is no spring constant, no mean-reversion, and no equilibrium position. The particle drifts indefinitely. This distinguishes the mapping from a harmonic oscillator, and the distinction matters: a harmonic oscillator has a stationary Boltzmann distribution, while a freely drifting Brownian particle does not. For the wealth distribution to have a steady state, additional ingredients (income, consumption, death) are required; see Section 8.

3.3 Summary of the mapping

The key identifications—wealth as e^x, expected return as drift velocity, volatility as noise strength, $D = \sigma^2/2$ as the Einstein relation—are collected in the "Dynamics" rows of Table 1. No approximation is involved: the two columns are the same mathematical object read in different notation.

3.4 Multiple assets and the portfolio

When the investor holds N assets with return vector $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_N)^\top$, volatilities $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_N)^\top$, and correlation matrix $\boldsymbol{\rho}$, the portfolio return is

$$\frac{\mathrm{d}W}{W} = \mathbf{w}^\top \boldsymbol{\mu}\, \mathrm{d}t + \mathbf{w}^\top \boldsymbol{\Sigma}^{1/2}\, \mathrm{d}\mathbf{B}_t\,, \tag{11}$$

where $\mathbf{w}$ is the vector of portfolio weights and $\boldsymbol{\Sigma} = \mathrm{diag}(\boldsymbol{\sigma})\, \boldsymbol{\rho}\, \mathrm{diag}(\boldsymbol{\sigma})$ is the covariance matrix. The log-wealth again follows a Langevin equation:

$$\mathrm{d}x = \left(\mathbf{w}^\top \boldsymbol{\mu} - \tfrac{1}{2}\mathbf{w}^\top \boldsymbol{\Sigma}\, \mathbf{w}\right) \mathrm{d}t \; + \; \sqrt{\mathbf{w}^\top \boldsymbol{\Sigma}\, \mathbf{w}}\; \mathrm{d}B_t\,, \tag{12}$$

where the scalar Brownian motion B_t drives the portfolio return and the effective diffusion coefficient is $D_{\mathrm{port}} = \frac{1}{2}\mathbf{w}^\top \boldsymbol{\Sigma}\, \mathbf{w}$.

Remark (Correlation structure)*.* The covariance matrix $\boldsymbol{\Sigma}$ plays the role of a coupling matrix in the multi-dimensional Langevin system. Its eigenvectors define the independent "modes" of the system: the largest eigenvalue corresponds to the aggregate market factor; smaller eigenvalues correspond to sector or idiosyncratic modes. This is mathematically identical to the normal-mode decomposition of coupled linear systems in physics, though the dynamics here are stochastic rather than deterministic. Random matrix theory (the Marčenko–Pastur law) provides the tools to separate signal from noise in the empirical covariance matrix—a problem with no counterpart in deterministic coupled oscillators.

4 The Fokker–Planck equation for wealth

4.1 From individual trajectories to distributions

Remark (Terminology)*.* Throughout this paper, we use *investor* for the individual decision-maker in the financial market, consistent with the asset pricing literature. In the statistical physics mapping, each investor's wealth trajectory corresponds to a *particle* trajectory in the Langevin/Fokker–Planck framework, and the population of investors corresponds to the *ensemble*. We reserve the term *agent* for the distinct context of agent-based computational models, which are not our primary framework here.

Suppose there are $\mathcal{N}$ investors, each with wealth $W_i(t)$ following (9) with common parameters (μ, σ). Here W denotes *market net wealth*: the market value of all assets minus the face value of all liabilities. When the wealth tax is levied at market value (the neutral case), the tax base equals W and the analysis is clean. When assets are assessed below market value while liabilities remain fully deductible, the tax base diverges from W—a distinction that becomes important in the distortion analysis of Section 7.1.

Let $p(W, t)$ denote the probability density: $p(W, t)\, \mathrm{d}W$ is the fraction of investors with wealth in $[W, W + \mathrm{d}W]$ at time t.

The standard result from stochastic calculus (the forward Kolmogorov equation) gives the

Fokker–Planck equation corresponding to the stochastic differential equation (9):

$$\boxed{\frac{\partial p}{\partial t} = -\frac{\partial}{\partial W}[\mu W\, p] + \frac{1}{2}\frac{\partial^2}{\partial W^2}[\sigma^2 W^2\, p]\,.} \tag{13}$$

The first term is the *drift* (or advection): wealth is pushed upward at rate μW. The second term is the *diffusion*: the distribution spreads due to the stochastic fluctuations of magnitude σW.

4.2 Fokker–Planck in log-wealth coordinates

Because the noise in (9) is multiplicative, the coefficients in (13) depend on W. The coordinate change $x = \ln W$ removes this state-dependence. Let $\pi(x,t)$ denote the density of log-wealth, related to $p(W,t)$ by

$$\pi(x,t) = W\, p(W,t) = e^x\, p(e^x,t)\,. \tag{14}$$

Substituting into (13), we obtain the Fokker–Planck equation for $\pi(x,t)$:

$$\boxed{\frac{\partial \pi}{\partial t} = -v\frac{\partial \pi}{\partial x} + D\frac{\partial^2 \pi}{\partial x^2}\,,} \tag{15}$$

where $v = \mu - \sigma^2/2$ is the drift velocity and $D = \sigma^2/2$ is the diffusion coefficient. This is the standard *drift–diffusion equation* with constant coefficients—the simplest non-trivial Fokker–Planck equation, and one of the most studied objects in statistical physics (Zwanzig, 2001, Ch. 3)(Livi and Politi, 2017, Ch. 6).

Remark (Propagator). The Green's function of (15) is the Gaussian

$$\pi(x,t\,|\,x_0,0) = \frac{1}{\sqrt{4\pi D t}}\exp\left(-\frac{(x-x_0-vt)^2}{4Dt}\right), \tag{16}$$

which is simply the statement that log-wealth is normally distributed with mean $x_0 + vt$ and variance $2Dt = \sigma^2 t$. This is the standard result that $W(t)$ is log-normally distributed under GBM. (The Gaussian form of the propagator is specific to the constant-coefficient case; the drift-shift symmetry that underlies our neutrality result does not depend on it—see Section 6.3.) The Fokker–Planck formulation adds nothing new for a single investor, but becomes essential when we consider distributions across investors and their evolution under policy changes.

4.3 Probability current

The Fokker–Planck equation can be written as a continuity equation:

$$\frac{\partial \pi}{\partial t} + \frac{\partial J}{\partial x} = 0\,, \tag{17}$$

where the *probability current* is

$$J(x,t) = v\,\pi(x,t) - D\frac{\partial \pi}{\partial x}\,. \tag{18}$$

The current J has a direct interpretation: it is the net flux of investors (in probability terms) flowing past the point x per unit time. The drift term $v\pi$ pushes probability to the right (higher wealth); the diffusion term $-D\,\partial\pi/\partial x$ drives probability from high-density to low-density regions (see Livi and Politi, 2017, Ch. 6, for boundary conditions and current analysis).

The concept of probability current will be central to the neutrality analysis: a tax is neutral if it does not alter the *relative* currents between any two points in wealth space.

5 Proportional wealth tax as drift modification

5.1 The taxed dynamics

Now impose a proportional wealth tax at rate τ_w on the market value of all holdings. Following Frøseth (2026b), the after-tax wealth dynamics are

$$\frac{\mathrm{d}W}{W} = (\mu - \tau_w)\,\mathrm{d}t + \sigma\,\mathrm{d}B_t\,. \tag{19}$$

The tax reduces the expected return from μ to $\mu - \tau_w$, but leaves the volatility σ unchanged. This is the multiplicative separability result: the tax operates as a proportional dilution, removing a fraction τ_w of wealth per unit time, without altering the stochastic structure of returns.

In log-wealth:

$$\mathrm{d}x = \underbrace{\left(\mu - \tau_w - \tfrac{\sigma^2}{2}\right)}_{\equiv\, v_\tau}\mathrm{d}t \;+\; \sigma\,\mathrm{d}B_t\,. \tag{20}$$

The only change is the drift velocity: $v \to v_\tau = v - \tau_w$.

5.2 The taxed Fokker–Planck equation

The Fokker–Planck equation for the taxed system is

$$\boxed{\frac{\partial\pi}{\partial t} = -v_\tau\frac{\partial\pi}{\partial x} + D\frac{\partial^2\pi}{\partial x^2}\,,} \tag{21}$$

with $v_\tau = v - \tau_w$ and $D = \sigma^2/2$ unchanged. The taxed probability current is

$$J_\tau(x,t) = v_\tau\,\pi(x,t) - D\frac{\partial\pi}{\partial x}\,. \tag{22}$$

5.3 What the tax does and does not change

Quantity	**Untaxed**	**Taxed**
Drift velocity	$v = \mu - \sigma^2/2$	$v_\tau = \mu - \tau_w - \sigma^2/2$
Diffusion coefficient	$D = \sigma^2/2$	$D = \sigma^2/2$
Noise strength	σ	σ
Propagator width	$\sqrt{2Dt}$	$\sqrt{2Dt}$
Propagator centre	$x_0 + vt$	$x_0 + v_\tau t$

The tax translates the centre of the propagator to the left (lower average wealth) but does not change its width (same dispersion around the mean). In physical terms: the particle drifts more slowly, but diffuses at the same rate. Figure 1 illustrates this for a population of investors evolving over twenty years.

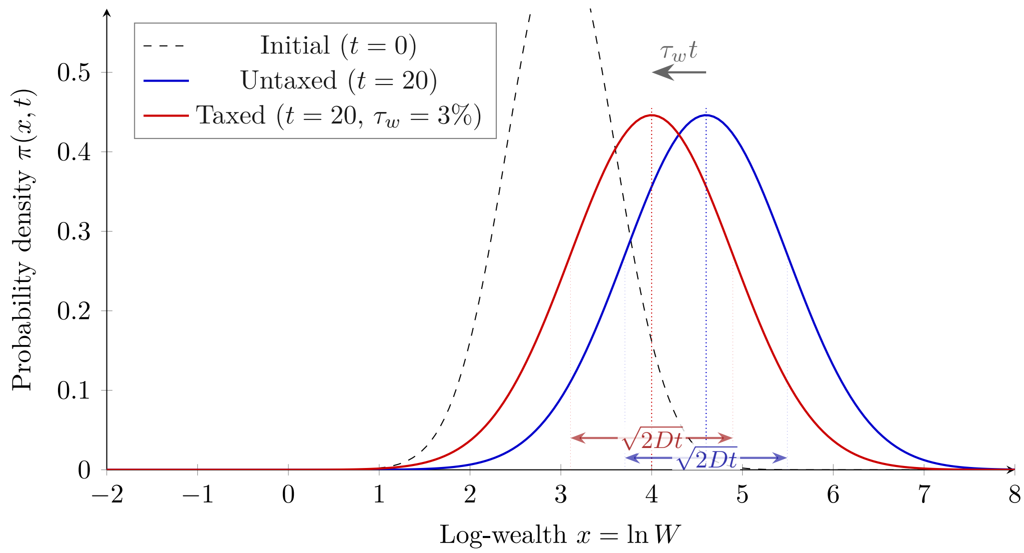

Figure 1: The separability result visualised. A population of investors starts at log-wealth $x_0 = 3$ (dashed). After $t = 20$ years with $\mu = 10\%$ and $\sigma = 20\%$, the untaxed distribution (blue) and the distribution under a $\tau_w = 3\%$ proportional wealth tax (red) have identical shape and width ($\sqrt{2Dt} = 0.89$) but different centres. The tax shifts the propagator to the left by $\tau_w t = 0.6$ in log-wealth without deforming it. This is the visual content of neutrality: the tax is a pure translation in log-wealth space.

6 Neutrality as a symmetry

6.1 The drift-shift symmetry

The central observation is that the proportional wealth tax acts as a *Galilean-type boost* in log-wealth space: it shifts the drift velocity uniformly, without coupling to the state, the diffusion, or any other parameter.

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

Proposition 2 (Neutrality as invariance)**.** *Let two assets have drift–diffusion parameters (v_1, D_1) and (v_2, D_2) under the untaxed Fokker–Planck equation. Under the proportional wealth tax at rate τ_w, the parameters become $(v_1 - \tau_w, D_1)$ and $(v_2 - \tau_w, D_2)$. Then:*

1. *The difference in drift velocities is unchanged: $(v_1 - \tau_w) - (v_2 - \tau_w) = v_1 - v_2$.*

2. *The ratio of diffusion coefficients is unchanged: D_1/D_2.*

3. *The Sharpe-ratio-like quantity $(v_i - v_j)/\sqrt{D_i + D_j - 2D_{ij}}$ is unchanged for all pairs (i, j).*

4. *The optimal portfolio weights, which depend only on $\boldsymbol{\mu} - r_f\mathbf{1}$ and $\boldsymbol{\Sigma}$, are unchanged because $(\boldsymbol{\mu} - \tau_w\mathbf{1}) - (r_f - \tau_w)\mathbf{1} = \boldsymbol{\mu} - r_f\mathbf{1}$.*

Proof. Properties (i)–(iii) follow directly from the linearity of $\mathcal{T}_\tau$ and the fact that it shifts all drifts by the same constant. Property (iv) uses the fact that the risk-free rate is also subject to the wealth tax: $r_f \to r_f - \tau_w$. The excess return $\mu_i - r_f$ is invariant. Since the Markowitz optimisation depends only on excess returns and the covariance matrix $\boldsymbol{\Sigma}$ (which involves only the diffusion coefficients), the optimal weights are unchanged. □

Remark (Physical interpretation). In the language of statistical physics, the tax is a uniform external field that couples identically to all degrees of freedom. Such a field shifts the equilibrium of every mode by the same amount; the relative structure—which modes are excited, which are suppressed—is preserved. This is precisely the content of the neutrality theorem in Frøseth (2026b): the tax contracts the opportunity set homothetically without distorting its shape.

6.2 Probability current and detailed balance

In an untaxed system with no income or consumption, the probability current is $J_0 = v\pi - D\,\partial\pi/\partial x$. In the taxed system, $J_\tau = (v - \tau_w)\pi - D\,\partial\pi/\partial x$. The *difference* is

$$J_\tau - J_0 = -\tau_w\,\pi(x, t)\,. \tag{23}$$

This is a uniform reduction of the probability current, proportional to the local density. No new currents are created between states that did not already have a current; no existing currents are reversed. The tax removes probability uniformly, like a spatially uniform decay rate.

If the untaxed system is in detailed balance ($J_0 = 0$ at steady state), then the taxed system has $J_\tau = -\tau_w\pi < 0$: a uniform leftward current representing the steady drain of wealth. Detailed balance is broken, but in the most benign way possible—the system is driven uniformly, with no state-dependent distortion. (Note that detailed balance requires a stationary distribution, which pure GBM does not possess; the confining mechanisms of Section 8 are needed. The argument here applies to the stationary regime of such an extended model.)

6.3 Robustness of the drift-shift symmetry

The preceding results were derived under geometric Brownian motion: normally distributed log-returns with constant drift and volatility. We now show that the drift-shift symmetry is robust to two important generalisations—non-Gaussian return distributions and stochastic volatility—because it rests on the *tax mechanism*, not on the distributional form.

6.3.1 Beyond Gaussian returns

The neutrality analysis in Frøseth (2026b) shows that the proportional wealth tax acts as a deterministic multiplicative scalar on wealth, independently of the return realisation. After n periods, a taxed investor's wealth is

$$W_n^{\text{taxed}} = (1 - \tau_w)^n \, W_n^{\text{untaxed}} \, . \tag{24}$$

This identity requires only two properties of the tax: proportionality (a constant fraction τ_w of all holdings) and universality (the same rate on every asset). It makes no reference to the distribution of returns. In log-wealth coordinates, the multiplicative factor becomes an additive constant:

$$x_n^{\text{taxed}} = x_n^{\text{untaxed}} + n \ln(1 - \tau_w) \, , \tag{25}$$

so the tax is a deterministic translation in log-wealth space, regardless of whether the increments of x_n^{untaxed} are Gaussian, Student-t, or drawn from any other distribution with well-defined moments.

Proposition 3 (Distribution-free drift shift)**.** *Let the return process be any continuous-time Itô diffusion*

$$\mathrm{d}x = \mu(x,t)\,\mathrm{d}t + \sigma(x,t)\,\mathrm{d}B_t \, , \tag{26}$$

with possibly state-dependent drift $\mu(x,t)$ and diffusion $\sigma(x,t)$. Under a proportional wealth tax at rate τ_w, the taxed log-wealth satisfies

$$\mathrm{d}x^{\text{taxed}} = [\mu(x,t) - \tau_w]\,\mathrm{d}t + \sigma(x,t)\,\mathrm{d}B_t \, . \tag{27}$$

The corresponding Fokker–Planck equation for the taxed density $\pi_\tau(x,t)$ is

$$\frac{\partial \pi_\tau}{\partial t} = -\frac{\partial}{\partial x}\left[(\mu(x,t) - \tau_w)\,\pi_\tau\right] + \frac{1}{2}\frac{\partial^2}{\partial x^2}\left[\sigma(x,t)^2\,\pi_\tau\right] . \tag{28}$$

The tax modifies only the drift coefficient $\mu \to \mu - \tau_w$; the diffusion coefficient $\sigma(x,t)^2/2$ is unchanged.

Proof. The proportional wealth tax removes a fraction τ_w of wealth per unit time. In the wealth-level dynamics $\mathrm{d}W = W[\mu_W(x,t)\,\mathrm{d}t + \sigma_W(x,t)\,\mathrm{d}B_t]$, the tax enters as an additional deterministic drain $-\tau_w W\,\mathrm{d}t$ in the drift, giving $\mathrm{d}W = W[(\mu_W - \tau_w)\,\mathrm{d}t + \sigma_W\,\mathrm{d}B_t]$. Applying Itô's lemma to $x = \ln W$ yields (27), since the Itô correction $-\sigma_W^2/2$ and the noise term $\sigma_W\,\mathrm{d}B_t$ are both independent of τ_w. The Fokker–Planck equation (28) follows from the standard forward

Kolmogorov equation for (27). □

Remark (Scope of the distribution-free result). Proposition 3 shows that the drift-shift structure $\mu \to \mu - \tau_w$ is preserved for any Itô diffusion, including processes with state-dependent drift and volatility (e.g. mean-reverting returns, local volatility models). The Gaussian propagator (16) is specific to the constant-coefficient case; the drift-shift symmetry is not. In particular, the neutrality result of Proposition 2—that relative drifts, diffusion ratios, and optimal portfolio weights are tax-invariant—extends to any return process of the form (26), provided the tax applies proportionally and universally.

The explicit formulas for the Pareto exponent (42) and the spectral gap (46) do use the constant-coefficient (GBM) assumption, because they depend on the specific form of the stationary distribution. The qualitative results—that the tax steepens the Pareto tail and that relaxation is slow—hold more broadly, but the exact functional forms would differ for non-constant coefficients.

6.3.2 Stochastic volatility

The GBM assumption also fixes the volatility σ as constant. Empirically, asset return volatility is time-varying, with well-documented clustering, mean reversion, and leverage effects (Heston, 1993). These features are captured by stochastic volatility models, in which the variance is itself a random process.

We consider the Heston model as a concrete and widely used example. The risky asset price and its instantaneous variance form a pair of coupled diffusions:

$$\frac{\mathrm{d}S}{S} = \mu\,\mathrm{d}t + \sqrt{v_t}\,\mathrm{d}W_t^{(1)}\,, \tag{29}$$

$$\mathrm{d}v_t = \lambda(\theta - v_t)\,\mathrm{d}t + \kappa\sqrt{v_t}\,\mathrm{d}W_t^{(2)}\,, \tag{30}$$

where $v_t = \sigma_t^2$ is the instantaneous variance, $\theta > 0$ the long-run mean, $\lambda > 0$ the rate of mean reversion, $\kappa > 0$ the volatility of variance, and $\mathrm{corr}(\mathrm{d}W^{(1)}, \mathrm{d}W^{(2)}) = \rho$. When $\rho < 0$, negative returns coincide with rising volatility (the leverage effect).

The wealth of an investor who allocates a fraction w to the risky asset and the remainder to a risk-free asset with rate r_f, subject to a proportional wealth tax, evolves as

$$\mathrm{d}W = W\left[r_f + w(\mu - r_f) - \tau_w\right]\mathrm{d}t + wW\sqrt{v_t}\,\mathrm{d}W_t^{(1)}\,. \tag{31}$$

The Fokker–Planck equation for the joint density $\pi(x, v, t)$ of log-wealth $x = \ln W$ and variance v is two-dimensional. Crucially, the tax rate τ_w appears only in the drift of x—the dynamics of v (Equation (30)) are entirely independent of the tax. This is the two-dimensional analogue of the drift-shift symmetry: the tax acts along the wealth axis only, leaving the volatility state untouched.

Proposition 4 (Portfolio neutrality under stochastic volatility)**.** *Under the Heston model* (29)–

(30) *with CRRA preferences* $U(C) = C^{1-\gamma}/(1-\gamma)$ *and a proportional wealth tax on all assets, the optimal portfolio weight* w^* *is independent of the wealth tax rate* τ_w*.*

Proof. Under CRRA utility, the value function admits the separable form $J(W, v, t) = W^{1-\gamma} f(v, t)/(1-\gamma)$, where $f > 0$ encodes the dependence on the volatility state and the investment horizon. The first-order condition for the portfolio weight in the Hamilton–Jacobi–Bellman equation yields (see Appendix B for the full derivation)

$$w^* = \underbrace{\frac{\mu - r_f}{\gamma v}}_{\text{myopic}} + \underbrace{\frac{f_v}{f} \cdot \frac{\kappa \rho}{\gamma}}_{\text{hedging}} . \tag{32}$$

The myopic demand depends on the excess return, risk aversion, and current variance—not on wealth or the tax rate. The hedging demand—the intertemporal component identified by Merton (1973)—depends on the ratio f_v/f, where $f(v, t)$ satisfies a PDE obtained by substituting the separable value function into the HJB equation. In this PDE, the tax rate τ_w appears only in a term $(1-\gamma)(r_f - \tau_w)f$ that shifts the effective discount rate but does not interact with v. Using the standard exponential-affine ansatz $f(v, t) = \exp(A(t) + B(t)v)$ (Chacko and Viceira, 2005), the Riccati equation for $B(t)$ collects only the v-dependent terms and is therefore independent of τ_w. Since $f_v/f = B(t)$, the hedging demand is tax-invariant. Consequently, both components of w^* are independent of τ_w. □

Corollary 1 (General Markov diffusions)**.** *The result extends to any Markov diffusion model with* K *risky assets and* M *state variables* $\mathbf{X}_t = (X_1, \ldots, X_M)^\top$*, in which the expected returns* $\mu_i(\mathbf{X})$*, volatilities* $\sigma_{ij}(\mathbf{X})$*, and state dynamics* $a_m(\mathbf{X})$*,* $b_{mj}(\mathbf{X})$ *depend on the state but not on the investor's wealth. Under CRRA preferences, the value function separates as* $J = W^{1-\gamma} f(\mathbf{X}, t)/(1-\gamma)$*, and the optimal portfolio weights*

$$\mathbf{w}^* = \frac{1}{\gamma}\mathbf{V}(\mathbf{X})^{-1}(\boldsymbol{\mu}(\mathbf{X}) - r_f\mathbf{1}) + \frac{1}{\gamma}\mathbf{V}(\mathbf{X})^{-1}\boldsymbol{\Phi}(\mathbf{X})\,\frac{\nabla_{\mathbf{X}} f}{f}$$

are independent of the wealth tax rate τ_w*. This encompasses the Heston, Hull–White, SABR, and affine term structure models as special cases (see Frøseth (2026a) for the full proof).*

Remark (Two mechanisms, one conclusion). The location-scale result (Section 6.3.1) and the stochastic volatility result rest on different mechanisms. The former uses the algebraic structure of the tax (multiplicative separability) and requires no utility specification. The latter requires CRRA preferences but places no restriction on the return distribution—Heston returns are not in the location-scale family. Together, they show that the drift-shift symmetry and its portfolio neutrality consequence are robust across a wide class of models. Appendix C develops a geometric interpretation that unifies both mechanisms in the language of fiber bundles and Galilean symmetry.

7 Distortion channels as symmetry breaking

The neutrality result of Section 6 rests on the wealth tax entering as a uniform, state-independent drift shift. Each of the distortion channels identified in Frøseth (2026a) breaks this condition in a specific way. We now classify them by the type of modification they introduce into the Fokker–Planck equation.

Remark (Physical intuition: gravity versus friction). The physical analogies are more than formal. The neutral wealth tax acts as a *uniform gravitational field* on the wealth distribution: it pulls every particle (investor) downward at the same rate, regardless of composition, without altering the thermal fluctuations (volatility). Gravity modifies the drift but not the diffusion—precisely the drift-shift symmetry. Because the tax payment $\tau_w W$ is deterministic given wealth, it enters the Langevin equation as a force, not as noise; a deterministic force can only modify the drift coefficient.

Each distortion channel introduces something beyond gravity. Liquidity frictions are literally *friction*: forced selling into illiquid markets dissipates wealth in a manner that depends on the state and couples to both drift and diffusion. Book-value assessment introduces an *anisotropic* field that pulls different assets at different rates. Migration creates an *absorbing boundary*—a cliff edge in the potential landscape. These analogies, summarised alongside the formal modifications below, may help readers carry a physical picture through the classification.

7.1 Channel 1: Book-value assessment

When assets are taxed at book value rather than market value, different assets attract different effective tax rates. If asset i has a book-to-market ratio β_i, the effective tax rate is $\tau_w^{(i)} = \tau_w \cdot \beta_i$. The drift shift becomes asset-dependent:

$$v_i \;\to\; v_i - \tau_w \beta_i \,. \tag{33}$$

The drift-shift transformation is no longer uniform: assets with low book-to-market ratios (growth stocks, intangible-heavy firms) are taxed less than assets with high book-to-market ratios (value stocks, asset-heavy firms). In the Fokker–Planck equation, the drift coefficient becomes state-dependent through the portfolio composition, since the portfolio-level effective tax rate depends on which assets the investor holds.

The leverage amplification. The anisotropic drift is amplified by the asymmetric treatment of debt in the net wealth tax base. Assets are assessed at book value ($\beta_i < 1$ for underassessed assets such as real estate), but liabilities are deducted at face value ($\beta_{\text{debt}} = 1$). An investor with assets of market value A assessed at βA and debt D has:

$$W_{\text{tax}} = \beta A - D \,, \qquad W_{\text{market}} = A - D \,, \tag{34}$$

so the ratio $W_{\text{tax}}/W_{\text{market}} = (\beta A - D)/(A - D)$ decreases with leverage and can become negative even when $W_{\text{market}} > 0$. This creates an incentive to lever up: borrowing against underassessed assets reduces the tax base disproportionately. In the Fokker–Planck equation, the effective drift depends not only on portfolio composition (through β_i) but also on the leverage ratio, coupling the debt decision to the wealth dynamics. Moreover, leveraged positions amplify the effective volatility of net wealth—connecting the assessment channel to the diffusion modification of Channel 2 below.

Symmetry broken: Uniformity of $\mathcal{T}_\tau$ across assets. The tax now couples differently to different degrees of freedom, distorting relative returns and portfolio choice. The asymmetric treatment of debt amplifies the distortion for leveraged investors.

7.2 Channel 2: Liquidity frictions

If the investor must sell assets to pay the tax and faces transaction costs or illiquidity, the effective tax rate depends on the liquidity of the portfolio. In the extreme, an investor holding only illiquid assets faces a higher effective burden than one holding liquid assets. The Fokker–Planck equation acquires a *state-dependent drift modification*:

$$v \;\to\; v - \tau_w - c(W, \ell)\,, \tag{35}$$

where $c(W, \ell)$ is a friction cost that depends on wealth W (or log-wealth x) and a liquidity parameter ℓ. For investors with wealth concentrated in illiquid assets, c is large; for liquid portfolios, $c \approx 0$.

Moreover, liquidity frictions can modify the *diffusion* coefficient if forced selling at unfavourable prices increases effective volatility:

$$D \;\to\; D + \Delta D(W, \ell)\,, \tag{36}$$

introducing a state-dependent diffusion that couples the tax to the stochastic structure of returns.

Symmetry broken: State-independence and drift-only coupling. The tax now modifies both drift and diffusion, and does so differently for different investors.

7.3 Channel 3: Forced dividend extraction

When firm owners extract dividends to pay the wealth tax, the firm's capital stock is reduced. If the firm faces financing frictions (limited credit, costly equity issuance), this extraction reduces investment and growth. The wealth tax then *couples the investor's dynamics to the firm's dynamics*:

$$\mu \;\to\; \mu(K) \quad \text{where} \quad \frac{\mathrm{d}K}{\mathrm{d}t} = f(K) - \delta K - \tau_w W\,, \tag{37}$$

and K is the firm's capital, $f(K)$ is the production function, δ is depreciation. The expected return μ is no longer a constant but depends on the firm's state, which in turn depends on the cumulative tax extracted. The Fokker–Planck equation acquires a *memory* through the coupling to the slow variable K.

Symmetry broken: Constancy of drift. The drift is endogenous, coupled to firm-level dynamics that evolve on a different timescale.

7.4 Channel 4: Migration

If investors can emigrate to escape the wealth tax, the system acquires an *absorbing boundary* in wealth space. An investor with wealth $W > W^*$ (where W^* is the threshold at which the tax burden exceeds the cost of emigration) may exit the system entirely. The Fokker–Planck equation is modified by a sink term:

$$\frac{\partial \pi}{\partial t} = -v_\tau \frac{\partial \pi}{\partial x} + D \frac{\partial^2 \pi}{\partial x^2} - \gamma(x)\, \pi\,, \tag{38}$$

where $\gamma(x)$ is a loss rate that increases sharply for $x > x^* = \ln W^*$. This drains probability from the upper tail of the distribution, depleting the wealthiest investors.

Symmetry broken: The system is no longer closed. The tax induces a state-dependent outflow that preferentially removes high-wealth investors, truncating the distribution.

7.5 Channel 5: Market impact

If forced asset sales to pay the tax move prices, the tax introduces a *flow-dependent* nonlinearity. Under the inelastic markets hypothesis (Gabaix and Koijen, 2022), aggregate selling pressure F depresses prices by an amount proportional to $\sqrt{F}$. The effective return becomes

$$\mu \;\to\; \mu - \alpha \sqrt{F(\tau_w, \mathcal{N})}\,, \tag{39}$$

where F depends on the tax rate and the number of investors selling simultaneously. This introduces a *collective effect*: the drift reduction experienced by each investor depends on the aggregate behaviour of all investors.

In the Fokker–Planck equation, this appears as a nonlinear, self-consistent drift:

$$v \;\to\; v - \tau_w - \alpha \sqrt{\tau_w \int W\, p(W,t)\, \mathrm{d}W}\,. \tag{40}$$

The drift depends on the first moment of the distribution itself, creating a feedback loop between the individual dynamics and the aggregate state.

Symmetry broken: Linearity and investor-independence. The tax couples each investor's dynamics to the aggregate distribution, introducing mean-field interactions absent from the untaxed system.

7.6 Summary of symmetry-breaking mechanisms

Channel	FP modification	Symmetry broken	Physics analogue
Neutral tax	$v \to v - \tau_w$	None	Uniform external field
1. Book value	$v_i \to v_i - \tau_w \beta_i$	Uniformity across assets	Anisotropic field
2. Liquidity	$v, D \to v(W), D(W)$	State-independence	State-dependent friction
3. Dividends	$\mu \to \mu(K(t))$	Drift constancy	Coupled slow variable
4. Migration	$+$ sink $\gamma(x)\pi$	Closed system	Absorbing boundary
5. Market impact	$v \to v(\langle W \rangle)$	Investor-independence	Mean-field interaction

Figure 2 illustrates how each channel deforms the propagator relative to the neutral (pure translation) case.

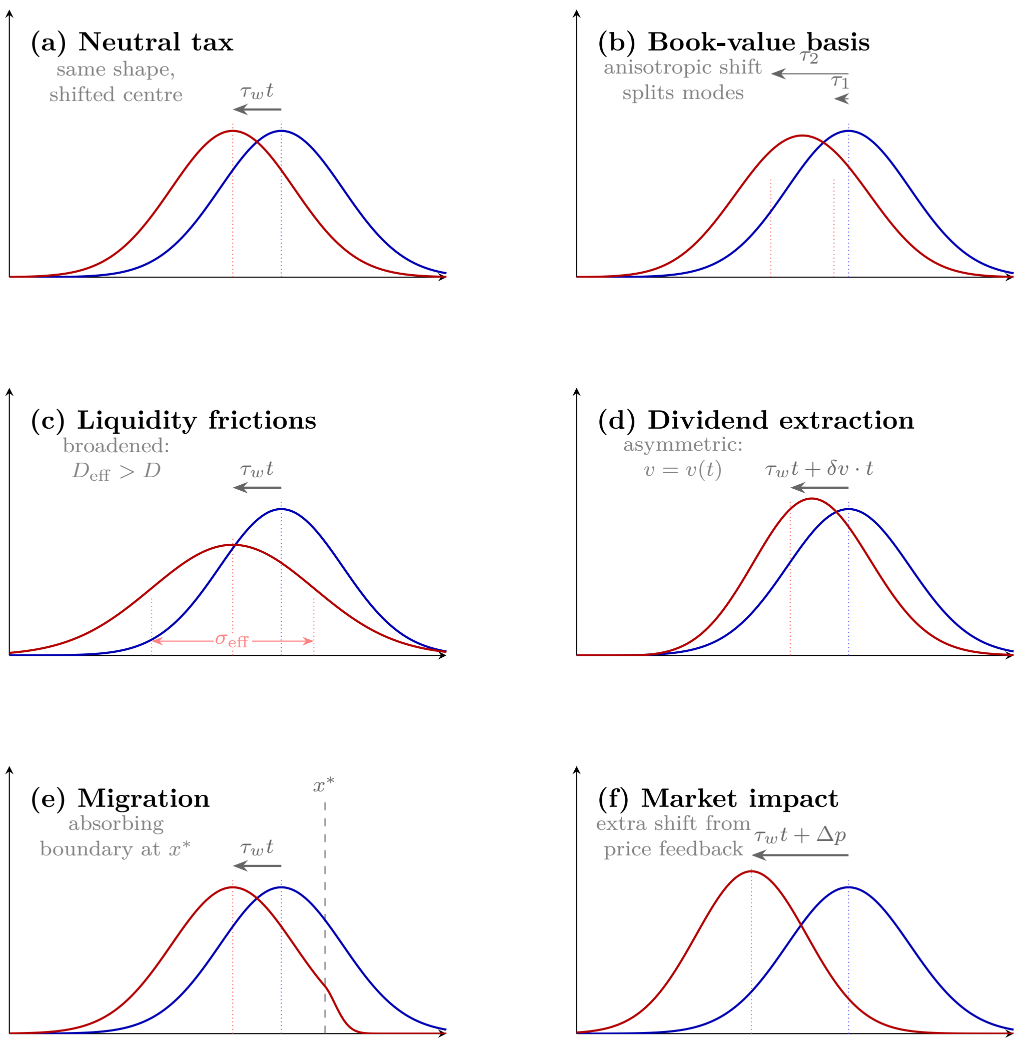

Figure 2: How each distortion channel deforms the wealth distribution relative to the neutral case. Blue: untaxed propagator. Red: taxed propagator. (a) Neutral tax: pure translation, no deformation. (b) Book-value assessment: different assets shift by different amounts, splitting the distribution into modes. (c) Liquidity frictions: the distribution broadens (increased effective D) as forced selling at unfavourable prices adds noise. (d) Dividend extraction: the drift becomes time-dependent as firm capital erodes, introducing asymmetry. (e) Migration: high-wealth investors exit the system above a threshold x^*, truncating the right tail. (f) Market impact: the collective selling pressure shifts the distribution further left than the tax alone would predict, with additional compression from price feedback.

8 Toward steady-state distributions

8.1 The problem with pure GBM

The Fokker–Planck equation (15) with constant v and D has no stationary solution on $(-\infty, \infty)$: the distribution drifts and spreads indefinitely. This reflects the well-known fact that geometric Brownian motion produces a log-normal distribution whose mean and variance both grow without bound.

For a wealth distribution to reach a steady state, the system must have sources and sinks that balance drift and diffusion.

8.2 Income and consumption as source and sink

In a realistic model, investors receive income and consume. The simplest extension adds an income flux λ (in units of log-wealth per time) and a consumption rate c, yielding

$$\mathrm{d}x = \left(v + \frac{\lambda}{W} - c\right)\mathrm{d}t + \sigma\,\mathrm{d}B_t\,, \tag{41}$$

where the λ/W term reflects that a fixed income λ has a larger effect on log-wealth when W is small. This creates a state-dependent drift that pushes low-wealth investors upward (income dominates) and allows high-wealth investors to drift further right (returns dominate). The resulting Fokker–Planck equation has a confining effect at low wealth and can admit a stationary solution.

8.3 Structure of the stationary distribution

Setting $\partial\pi/\partial t = 0$ in the modified Fokker–Planck equation with income and consumption yields a second-order ODE for $\pi_{\text{ss}}(x)$. While the general solution depends on the specific functional forms, the qualitative structure is well established in the econophysics literature (Yakovenko and Rosser, 2009):

- **Bulk** (low to moderate wealth): Approximately exponential (Boltzmann–Gibbs) distribution, $\pi \propto e^{-x/T}$, where T is an effective temperature related to the average wealth.
- **Tail** (high wealth): Power-law (Pareto) distribution, $p(W) \propto W^{-\alpha}$, arising from the multiplicative nature of returns at high wealth where income is negligible relative to capital gains.

The wealth tax modifies the Pareto exponent. When GBM is supplemented by an additive income or redistribution term—turning the dynamics into a Kesten process—the stationary distribution develops a Pareto tail (Kesten, 1973; Bouchaud and Mézard, 2000); see Gabaix (2009) for a survey of the random-growth mechanism and its applications. The tail exponent satisfies (see Appendix A for the derivation)

$$\alpha(\tau_w) = 1 - \frac{2(\mu - \tau_w)}{\sigma^2}\,. \tag{42}$$

The condition $\alpha > 0$ requires $\mu - \tau_w < \sigma^2/2$, i.e. the log-wealth drift must be negative for a stationary distribution to exist. Increasing the tax rate increases α, making the tail steeper (less inequality): the tax reduces the effective drift that sustains the power law, compressing the upper tail of the wealth distribution.

Remark (Neutrality and the distribution). There is no contradiction between the neutrality result (individual portfolio choice unchanged) and the distributional effect (Pareto exponent changes). Neutrality is a statement about each investor's optimisation problem; the distributional effect is a statement about the ensemble. The tax is neutral in the sense that no investor wants to change their portfolio, but the aggregate distribution shifts because all investors are uniformly poorer.

Remark (Distinction from Bouchaud–Mézard redistribution). The proportional wealth tax studied here is structurally different from the redistribution mechanism in Bouchaud and Mézard (2000). Their model features a mean-field exchange term $J \sum_j (w_j - w_i)$ that transfers wealth between agents—an off-diagonal coupling in the multi-agent Langevin system. The proportional wealth tax, by contrast, modifies each investor's own drift uniformly: $v \to v - \tau_w$, a diagonal perturbation that leaves the diffusion and all inter-investor couplings unchanged. These are distinct modifications of the Fokker–Planck equation. The exchange mechanism confines the distribution through inter-agent interactions; the wealth tax confines it through a shift in each investor's individual growth rate. The drift-shift symmetry that underlies our neutrality result has no analogue in the exchange framework.

8.4 Ergodicity and the Pareto condition

The Pareto exponent (42) admits a revealing reinterpretation through the lens of ergodicity (Peters, 2019). Under geometric Brownian motion, the ensemble-average growth rate of wealth $\mathbb{E}[\mathrm{d}W/W]/\mathrm{d}t = \mu$ differs from the time-average (almost-sure) growth rate of log-wealth,

$$g_{\text{time}} = \mu - \frac{\sigma^2}{2}, \tag{43}$$

whenever $\sigma > 0$. The gap between them is $\sigma^2/2 = D$, exactly the diffusion coefficient in the Fokker–Planck equation. This is the signature of non-ergodicity for multiplicative processes: the typical trajectory and the ensemble average diverge, and the divergence is governed by the same parameter that controls the spreading of the wealth distribution (see Zwanzig, 2001, for ergodicity in Brownian systems)(and Bouchaud and Farmer, 2021, for quasi-non-ergodicity in wealth dynamics).

The wealth tax reduces both averages uniformly: the ensemble-average growth rate becomes $\mu - \tau_w$ and the time-average becomes

$$v_\tau \;=\; \mu - \tau_w - \frac{\sigma^2}{2}, \tag{44}$$

which is the drift of taxed log-wealth that appears throughout the preceding sections. The Pareto exponent can now be written directly in terms of this time-average growth rate:

$$\alpha(\tau_w) = -\frac{2\,v_\tau}{\sigma^2} = -\frac{v_\tau}{D}. \tag{45}$$

This is a drift-to-diffusion ratio—the analogue of a Péclet number in transport physics—measuring the strength of directed motion relative to diffusive spreading. A physicist will recognise it immediately: the Pareto tail steepness is set by the competition between drift (which concentrates wealth) and diffusion (which disperses it).

Three consequences follow.

First, *the condition for a stationary distribution ($\alpha > 0$) is equivalent to $v_\tau < 0$*: the time-average growth rate of taxed log-wealth must be negative. This is not a pathology. It means the typical investor's wealth is shrinking over time, even as the ensemble mean grows (because the ensemble average is dominated by a few lucky trajectories in the upper tail). The tension between multiplicative growth for the fortunate few and mean-reversion via the additive Kesten term is precisely what creates and sustains the Pareto tail.

Second, *the wealth tax steepens the Pareto tail by making v_τ more negative.* This is the distributional counterpart of the neutrality result: the tax does not change any investor's portfolio (neutrality), but by shifting v_τ downward it tips the drift–diffusion balance further toward diffusion dominance, compressing the upper tail of the wealth distribution.

Third, *the non-ergodicity gap $\sigma^2/2 = D$ is invariant under the wealth tax.* Since the tax modifies only the drift, not the diffusion coefficient, it cannot alter the fundamental non-ergodic character of the dynamics. The gap between what the ensemble predicts and what the typical investor experiences remains unchanged. This invariance will reappear in the analysis of relaxation times, where D sets the timescale over which the distribution approaches its new steady state.

Finally, the neutrality result acquires a clean ergodic restatement. Consider a portfolio of N assets with time-average growth rates $g_i = \mu_i - \sigma_i^2/2$. The wealth tax transforms each rate to $g_i - \tau_w$, preserving the ranking: if $g_i > g_j$ before tax, then $g_i - \tau_w > g_j - \tau_w$ after. Since long-run portfolio choice depends on the ranking of time-average growth rates (the asset that compounds fastest almost surely dominates), the tax is neutral for portfolio selection.

8.5 Relaxation dynamics and the spectral gap

The ergodic reinterpretation tells us *where* the wealth distribution converges to—a Pareto tail with exponent α. It does not tell us *how fast*. After a tax change shifts the drift, the distribution relaxes toward a new steady state. The speed of this relaxation is governed by the *spectral gap* of the Fokker–Planck operator: the magnitude of the second eigenvalue of the Kolmogorov forward operator $\mathcal{A}^*$ (for the general eigenvalue theory, see Livi and Politi, 2017, Ch. 6).

Gabaix et al. (2016) prove that for the random growth process $\mathrm{d}x = \mu\,\mathrm{d}t + \sigma\,\mathrm{d}B_t$ with reflecting barrier at $x = 0$ and demographic turnover at rate δ (investors die and are replaced at a reference level), the cross-sectional distribution $p(x,t)$ converges exponentially to its stationary distribution $p_\infty(x)$:

$$\|p(\cdot,t) - p_\infty(\cdot)\| \;\sim\; k\,e^{-\lambda t}\,,$$

where $\|\cdot\|$ is the L^1 (total variation) norm and the rate of convergence is

$$\lambda \;=\; \frac{v_\tau^2}{2\sigma^2}\,\mathbf{1}_{\{v_\tau<0\}} \;+\; \delta\,. \tag{46}$$

The indicator function $\mathbf{1}_{\{v_\tau<0\}}$ is the key: the drift-dependent term contributes to convergence *only* when the time-average growth rate is negative ($v_\tau < 0$), so that the drift pushes probability back toward the barrier. When $v_\tau > 0$, the drift carries wealth away from the barrier, and only demographic turnover δ drives convergence. A physicist will recognise two distinct relaxation mechanisms: confining drift (analogous to a restoring force) and population renewal (analogous to coupling to a thermal reservoir).

The corresponding half-life is

$$t_{1/2} \;=\; \frac{\ln 2}{\lambda}\,. \tag{47}$$

Three features of (46) deserve emphasis.

First, *the Pareto exponent and the relaxation rate are linked.* When $v_\tau < 0$, the Pareto exponent from Section 8.4 is $\alpha = -2v_\tau/\sigma^2$, so the drift-dependent contribution to λ can be written

$$\lambda \;=\; \frac{\alpha^2\sigma^2}{8} + \delta\,. \tag{48}$$

The two key observables of the wealth distribution—the tail steepness (α) and the convergence rate (λ)—are related through a single formula. A steeper tail (larger α) implies faster convergence, because a stronger confining drift (more negative v_τ) both steepens the Pareto tail and accelerates the return to equilibrium.

Second, *there is a critical threshold at* $v_\tau = 0$ (equivalently, $\tau_w = \mu - \sigma^2/2$, or $\alpha = 0$). Below this threshold, the drift-dependent mechanism is inactive: convergence relies entirely on demographic turnover at rate δ. Above it, both mechanisms operate. This is the relaxation-dynamics counterpart of the steady-state result that the Pareto tail exists only for $\alpha > 0$.

Third, *for realistic parameters, relaxation is slow.* With $\sigma = 0.30$ (cross-sectional wealth volatility), $\mu = 0.08$ (expected return), $\delta = 1/30$ (generational turnover), and a tax rate $\tau_w = 2\%$: the taxed drift is $v_\tau = 0.08 - 0.02 - 0.045 = 0.015 > 0$. Since v_τ remains positive, the drift-dependent term is inactive, giving $\lambda = \delta = 0.033$ and $t_{1/2} \approx 21$ years. Even at $\tau_w = 5\%$, $v_\tau = -0.015$ and the drift contribution is only $0.015^2/(2 \times 0.09) \approx 0.001$, barely altering the half-life. Only at much higher tax rates (or lower expected returns) does the drift mechanism materially accelerate convergence.

This is the wealth-tax application of Gabaix et al.'s central finding: the baseline random growth model generates inherently slow transition dynamics. The wealth tax changes *where* the distribution converges to—the steady-state Pareto exponent α responds immediately to the new drift—but it has much less effect on *how fast* it gets there. Short-run distributional effects of a tax change may therefore look very different from the long-run steady state, with the transition stretching over decades. Figure 3 illustrates this for three tax rates.

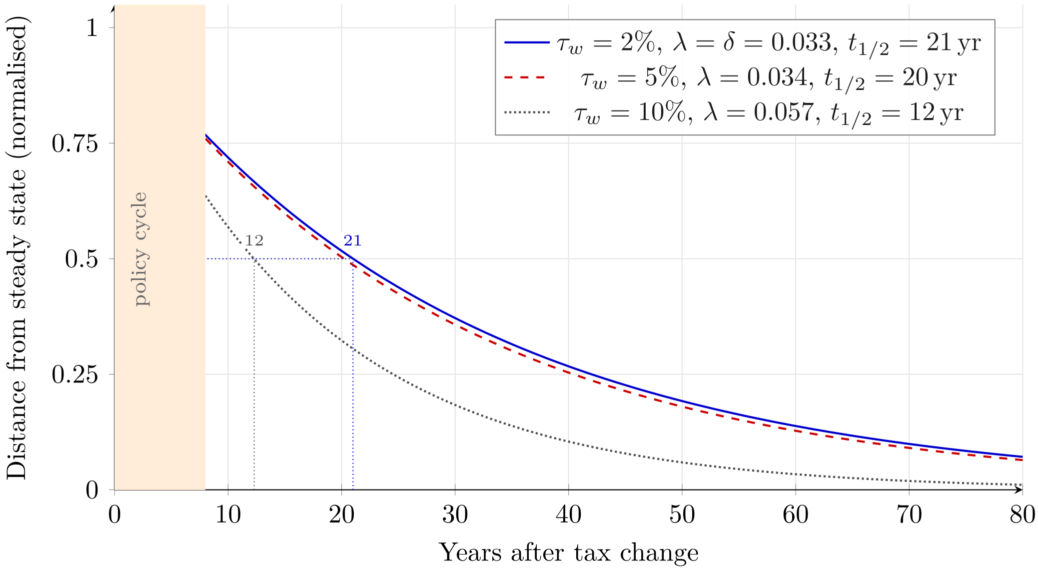

Figure 3: Relaxation toward the new steady-state wealth distribution after a tax change, for three tax rates. Parameters: $\sigma = 0.30$, $\mu = 0.08$, $\delta = 1/30$ (generational turnover). The vertical axis measures the distance $\|p(\cdot, t) - p_\infty\|$ normalised to unity at $t = 0$. At $\tau_w = 2\%$ the time-average growth rate remains positive ($v_\tau > 0$), so convergence depends entirely on demographic turnover ($\lambda = \delta$). At $\tau_w = 5\%$, v_τ turns negative but only marginally, adding less than 4% to the convergence rate. Even at the extreme rate $\tau_w = 10\%$, the half-life is still over a decade. The shaded region marks a typical electoral cycle (∼4–8 years), during which the distribution has moved less than a quarter of the way to its new steady state in all three scenarios.

9 Discussion and extensions

9.1 What the mapping enables

The preceding sections established six results: neutrality is a drift-shift symmetry (Section 6) that is robust to non-Gaussian returns and stochastic volatility (Section 6.3), each distortion channel breaks it in a classifiable way (Section 7), the Fokker–Planck equation governs the distributional consequences (Section 8), the Pareto tail steepness is a drift-to-diffusion ratio whose sign is controlled by the time-average growth rate (Section 8.4), and the spectral gap of the Fokker–Planck operator governs how fast the distribution converges to steady state (Section 8.5). Together, these open concrete avenues that the standard finance analysis does not provide:

1. **Empirical tests of symmetry breaking.** Each distortion channel predicts a specific deformation of the propagator (Figure 2). In principle, panel data on wealth trajectories before and after a tax change could be tested against these signatures—distinguishing, for example, broadening (liquidity) from truncation (migration).

2. **Quantitative bounds on distortion magnitude.** Because each channel modifies a specific coefficient in the Fokker–Planck equation, its effect can be bounded by estimating the relevant parameter (book-to-market dispersion, bid–ask spreads, migration elasticities). The taxonomy converts a qualitative policy debate into a parameter-estimation problem.

3. **Distributional dynamics.** The Fokker–Planck equation describes not only the steady-

state distribution but also the transient path toward it. The spectral gap analysis of Section 8.5 shows that this path can be slow—decades for realistic parameters—implying that short-run distributional effects of a tax change may differ substantially from the long-run steady state. These questions are inaccessible from within representative-agent asset pricing.

4. **Policy timescales.** Two results of the preceding analysis combine into a concrete policy message. First, the ergodicity analysis (Section 8.4) shows that a stationary wealth distribution *requires* the typical investor's log-wealth to shrink over time ($v_\tau < 0$). In real economies, median wealth growth is indeed lower than mean wealth growth—this is not a pathology of the model but the standard mechanism by which multiplicative noise generates heavy tails. The wealth tax steepens the Pareto tail by pushing the time-average growth rate further into negative territory, but the effect on the *typical* investor is a uniform reduction in growth, not a reallocation across assets. Second, the spectral gap (Section 8.5) quantifies *how long* these distributional shifts take. For realistic calibrations (Figure 3), the half-life of the transition exceeds twenty years—far longer than a typical electoral or policy cycle. A government introducing a wealth tax to reduce inequality should therefore expect a slow, multi-generational adjustment. Short-run revenue effects will materialise immediately (they depend only on the current distribution and the tax rate), but the distributional steady state lies decades away. This mismatch between fiscal timescales and distributional timescales is invisible in static or representative-agent analyses; it emerges naturally from the Fokker–Planck framework.

9.2 Open questions

Several extensions merit investigation.

Detailed balance, market efficiency, and nonlinear Fokker–Planck dynamics. The drift-shift symmetry of Section 6 has a natural equilibrium interpretation. Under the efficient market hypothesis, risk-adjusted returns are unpredictable: the drift under the risk-neutral measure equals the risk-free rate. In Fokker–Planck language, this is a detailed-balance condition—the probability current J vanishes at steady state (cf. Livi and Politi, 2017, Ch. 6). The proportional wealth tax preserves detailed balance because it shifts the drift uniformly, which is another way to see why prices are unaffected. The inelastic markets hypothesis of Gabaix and Koijen (2022) breaks detailed balance by introducing flow-dependent pricing—a mechanism given a microstructural interpretation by Bouchaud (2022)—so that the drift becomes a functional of the distribution itself. Under inelastic markets, the wealth tax creates a permanent probability current that has no counterpart in the efficient-markets case, and this current drives a persistent price impact. Modelling this requires a nonlinear Fokker–Planck equation of the form in (40), where the drift is self-consistently determined by the distribution. This moves the analysis from the linear, exactly solvable regime of the present paper into the territory of nonlinear diffusion and mean-field models—a substantial extension that we defer to future work.

Heterogeneous growth regimes. The present paper assumes a common drift μ and volatility σ for all investors. In practice, different wealth levels may face systematically different growth rates—for instance, if wealthy investors have access to higher-return asset classes. Gabaix et al. (2016) show that introducing wealth-dependent growth rates can dramatically accelerate the convergence to the stationary distribution, resolving the "slow transition" puzzle that arises when the spectral gap is dominated by the demographic term δ (Section 8.5). In the Fokker–Planck framework, this corresponds to a state-dependent drift $v(x)$, which breaks the constant-coefficient structure that underlies the exact results of this paper. Whether the neutrality result survives approximately in such a setting—and if so, under what conditions—is an open question with direct empirical relevance.

Debt, leverage, and the net wealth tax base. The leverage amplification identified in Section 7.1 implies that the Fokker–Planck dynamics of *taxable* net wealth can differ substantially from those of market net wealth. When assets are assessed below market value but debt is fully deductible, the effective drift and diffusion of the taxable distribution depend on the leverage ratio, which is itself an endogenous choice variable. A full treatment would model the joint dynamics of assets, debt, and taxable net wealth as a multi-dimensional Fokker–Planck system, with the leverage decision coupling the dimensions. This is particularly relevant for real estate, where assessment discounts are large and mortgage financing is prevalent. We defer the formal multi-dimensional treatment to a separate paper on redistribution analysis (Frøseth, 2026c), which provides a framework for incorporating it through the taxonomy of Fokker–Planck modifications.

Phase transitions. In principle, a sufficiently large change in the drift coefficient could trigger a qualitative change in the distribution—for example, a transition from a power-law tail to an exponential tail, or the emergence of bimodality. Whether such transitions occur at realistic tax rates is an open empirical and theoretical question.

A Derivation of the Pareto exponent

The formula for $\alpha(\tau_w)$ in (42) does not follow from pure GBM, which has no stationary distribution (Section 8). It arises when multiplicative growth is supplemented by an additive component—income, redistribution, or a reflecting boundary—so that the process takes the Kesten form

$$W_{n+1} = A_n W_n + B_n \,, \tag{49}$$

where A_n (the multiplicative factor from returns and tax) and B_n (the additive income/redistribution term) are i.i.d. and independent of each other. The classical result of Kesten (1973) states that if $\mathbb{E}[\ln A] < 0$ (so that wealth does not diverge) and B has suitable integrability, the stationary distribution has a Pareto tail: $\Pr(W > w) \sim w^{-\alpha}$, where $\alpha > 0$ is the unique solution of

$$\mathbb{E}\big[A^{\alpha}\big] = 1 \,. \tag{50}$$

In continuous time, the taxed GBM over an interval $\mathrm{d}t$ gives the multiplicative factor

$$A = \exp\Big[\big(\mu - \tau_w - \tfrac{\sigma^2}{2}\big)\mathrm{d}t + \sigma\sqrt{\mathrm{d}t}\, Z\Big], \qquad Z \sim \mathcal{N}(0,1) \,.$$

Evaluating the moment-generating function:

$$\mathbb{E}[A^\alpha] = \exp\Big[\alpha(\mu - \tau_w - \tfrac{\sigma^2}{2})\mathrm{d}t + \tfrac{\alpha^2\sigma^2}{2}\mathrm{d}t\Big].$$

Setting $\mathbb{E}[A^\alpha] = 1$ requires the exponent to vanish:

$$\alpha(\mu - \tau_w - \tfrac{\sigma^2}{2}) + \frac{\alpha^2\sigma^2}{2} = 0\,.$$

Dividing by $\alpha \neq 0$ and solving:

$$\alpha = -\frac{2(\mu - \tau_w - \sigma^2/2)}{\sigma^2} = \frac{\sigma^2 - 2(\mu - \tau_w)}{\sigma^2} = 1 - \frac{2(\mu - \tau_w)}{\sigma^2}\,,$$

which is (42). The condition $\alpha > 0$ requires $\mu - \tau_w < \sigma^2/2$, i.e. the log-wealth drift $v_\tau = \mu - \tau_w - \sigma^2/2$ must be negative—the multiplicative dynamics must be mean-reverting on average for a stationary distribution to exist.

B Portfolio neutrality under stochastic volatility

This appendix provides the full derivation of Proposition 4, establishing that the optimal portfolio weight under the Heston stochastic volatility model is independent of the wealth tax rate.

Setup

The investor maximises expected discounted CRRA utility $\mathbb{E}\Big[\int_0^T e^{-\delta t}\, C_t^{1-\gamma}/(1-\gamma)\,\mathrm{d}t\Big]$ subject to the wealth dynamics (31) and the variance dynamics (30). The value function is

$$J(W, v, t) = \max_{C,w}\ \mathbb{E}_t\left[\int_t^T e^{-\delta(s-t)}\frac{C_s^{1-\gamma}}{1-\gamma}\,\mathrm{d}s\right].$$

HJB equation

By Bellman's principle, J satisfies

$$\begin{aligned} 0 = \max_{C,w}\Big\{ & \frac{C^{1-\gamma}}{1-\gamma} - \delta J + J_t \\ & + J_W\big[W(r_f + w(\mu - r_f) - \tau_w) - C\big] \\ & + J_v\,\lambda(\theta - v) \\ & + \tfrac{1}{2}J_{WW}\,w^2 v W^2 + \tfrac{1}{2}J_{vv}\,\kappa^2 v \\ & + J_{Wv}\,wW\kappa v\rho\Big\}. \end{aligned} \tag{51}$$

Separable value function

The homogeneity of CRRA utility suggests the ansatz

$$J(W, v, t) = \frac{W^{1-\gamma}}{1-\gamma} f(v, t) , \tag{52}$$

with $f > 0$. The relevant partial derivatives are

$$\begin{aligned} J_W &= W^{-\gamma} f , & J_{WW} &= -\gamma W^{-\gamma-1} f , \\ J_v &= \frac{W^{1-\gamma}}{1-\gamma} f_v , & J_{vv} &= \frac{W^{1-\gamma}}{1-\gamma} f_{vv} , \\ J_{Wv} &= W^{-\gamma} f_v . \end{aligned} \tag{53}$$

First-order condition for w

Differentiating (51) with respect to w and setting the result to zero:

$$J_W W(\mu - r_f) + J_{WW}\, wvW^2 + J_{Wv}\, W\kappa v\rho = 0 .$$

Substituting (53) and dividing by $W^{1-\gamma} f$:

$$(\mu - r_f) - \gamma\, w\, v + \frac{f_v}{f} \kappa v\rho = 0 .$$

Solving for w^*:

$$w^* = \frac{\mu - r_f}{\gamma v} + \frac{f_v}{f} \cdot \frac{\kappa\rho}{\gamma} . \tag{54}$$

The tax rate τ_w does not appear in (54).

PDE for $f(v, t)$ and the Riccati argument

Substituting the separable form (52) and the optimal controls into the HJB equation, and dividing by $W^{1-\gamma}/(1-\gamma)$, yields a PDE for f:

$$\begin{aligned} 0 = {} & f_t + \lambda(\theta - v) f_v + \tfrac{1}{2}\kappa^2 v\, f_{vv} \\ & + h(v; \gamma, \mu, r_f, \kappa, \rho)\, f \\ & + (1-\gamma)(r_f - \tau_w)\, f + g(f) , \end{aligned} \tag{55}$$

where $h(v; \cdot\,)$ collects the v-dependent terms arising from the optimal portfolio (a quadratic in f_v/f) and $g(f)$ collects the contribution from optimal consumption. The key observation: τ_w appears *only* in the term $(1-\gamma)(r_f - \tau_w)f$, which is independent of v.

Using the exponential-affine ansatz $f(v, t) = \exp(A(t) + B(t)\, v)$ (Chacko and Viceira, 2005), substitution into (55) separates into:

- A **Riccati equation for** $B(t)$, arising from the terms proportional to v:

$$\dot{B} = -\lambda B + \tfrac{1}{2}\kappa^2 B^2 + h_1(\gamma, \mu, r_f, \kappa, \rho)\,,$$

  where h_1 absorbs the v-coefficient from h. This equation involves only the parameters of the variance dynamics and the risk-return trade-off—*not* τ_w.

- An **ODE for** $A(t)$, arising from the constant terms:

$$\dot{A} = \lambda\theta\, B + (1-\gamma)(r_f - \tau_w) + h_0\,,$$

  which absorbs τ_w alongside other constants.

Since $f_v/f = B(t)$ and $B(t)$ is independent of τ_w, the hedging demand $\frac{f_v}{f} \cdot \frac{\kappa\rho}{\gamma} = B(t) \cdot \frac{\kappa\rho}{\gamma}$ is tax-invariant. Combined with the tax-invariance of the myopic demand, this establishes that w^* is independent of τ_w. □

C Geometric interpretation of neutrality

The drift-shift symmetry and the CRRA separability result admit a geometric formulation that unifies the algebraic arguments of Section 6.3 and connects them to structures familiar from theoretical physics. The formulation uses the language of fiber bundles and connections (see Nakahara, 2003, Ch. 9–10, for a physicist's introduction); readers interested in the application of gauge-theoretic ideas to finance may consult Ilinski (2001).

The state space as a fiber bundle

Consider the investor's full state as a point in a product space: the *base* is the log-wealth coordinate $x \in \mathbb{R}$ (or, under stochastic volatility, the pair $(x, v) \in \mathbb{R} \times \mathbb{R}_+$), and the *fiber* above each base point is the portfolio simplex $\Delta^{N-1} = \{\mathbf{w} \in \mathbb{R}^N : \sum_i w_i = 1\}$. The investor's optimisation problem selects a *section* of this bundle: a rule $\mathbf{w}^*(x, \mathbf{X}, t)$ that assigns an optimal portfolio to each state.

Under CRRA utility, the value function separates as $J = W^{1-\gamma} f(\mathbf{X}, t)/(1-\gamma)$, and the first-order conditions yield a portfolio $\mathbf{w}^*$ that is independent of the base coordinate x (equivalently, of W). In differential geometry language, the optimal section is *horizontal*: it does not vary along the base. The connection on the bundle is *flat*—there is no curvature, no holonomy, and no wealth-dependence in the portfolio prescription.

The proportional wealth tax acts as a *vertical automorphism*: it translates every point along the base ($x \to x - \tau_w t$) without rotating the fiber. Because the connection is flat and the tax acts only vertically, the horizontal section is invariant. This is the geometric content of portfolio neutrality. Figure 4 illustrates the construction.

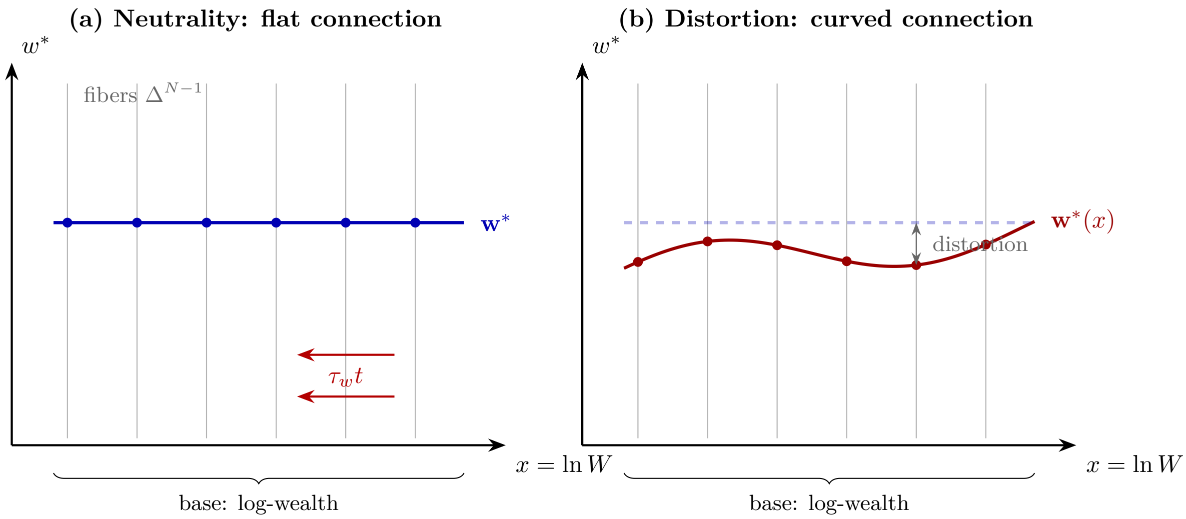

Figure 4: Geometric interpretation of neutrality and its breakdown. **(a)** Under CRRA preferences, the optimal portfolio $\mathbf{w}^*$ is a horizontal section of the fiber bundle: it does not vary with log-wealth x. The proportional wealth tax (red arrows) translates the base leftward without rotating the fibers, leaving the section invariant. **(b)** When a distortion channel is active (e.g. book-value assessment, liquidity frictions), the section becomes wealth-dependent: the connection acquires curvature, and the portfolio varies with x. The dashed line shows the flat (neutral) section for comparison.

Galilean symmetry in log-wealth space

The drift-shift transformation $\mathcal{T}_\tau : v \mapsto v - \tau_w$, $D \mapsto D$ (Section 6) is the analogue of a *Galilean boost* in classical mechanics (for the symmetry group of the diffusion equation, see Risken, 1989, Ch. 6). In the Galilean group, a uniform velocity shift preserves all relative velocities and all forces (which depend on position differences). Here, the tax shifts all drift velocities by the same constant τ_w, preserving all relative drifts, diffusion coefficients, and Sharpe ratios.

The portfolio weights, being functions of excess returns $\mu_i - r_f$ and the covariance matrix $\mathbf{\Sigma}$, are *Galilean invariants*: quantities that depend on velocity differences, not absolute velocities. The risk-free rate r_f transforms under the same boost ($r_f \to r_f - \tau_w$), so the excess return $\mu_i - r_f$ is invariant—exactly as relative velocities are invariant under a Galilean transformation.

This analogy makes precise the paper's characterisation of the tax as a "uniform external field" (Section 6): a spatially uniform field shifts the equilibrium of every degree of freedom by the same amount, preserving the relative structure.

Noether's theorem and the conserved charge

The one-parameter family of transformations $\{\mathcal{T}_\tau\}_{\tau \geq 0}$ is a continuous symmetry group of the portfolio optimality conditions. Noether's theorem associates a conserved quantity to each continuous symmetry. The conserved "charge" here is the excess return vector

$$\mathbf{q} \;\equiv\; \boldsymbol{\mu} - r_f\,\mathbf{1}\,, \qquad (56)$$

which is invariant under $\mathcal{T}_\tau$ because $(\boldsymbol{\mu} - \tau_w \mathbf{1}) - (r_f - \tau_w)\mathbf{1} = \boldsymbol{\mu} - r_f \mathbf{1}$. The optimal portfolio depends only on this conserved charge and the covariance matrix (which lives entirely in the diffusion sector and is therefore τ_w-independent). Portfolio neutrality is the statement that the optimal portfolio is a function of conserved quantities only.

Symmetry breaking as curvature

Each distortion channel of Section 7 has a geometric reading: it introduces *curvature* into the previously flat connection, making the optimal portfolio wealth-dependent. Book-value assessment creates a twist between the base and the fiber (the effective tax rate depends on unrealised gains, which depend on x). Liquidity frictions make the connection density-dependent. Migration introduces a boundary that breaks the translation symmetry of the base space. In each case, the horizontal section ceases to exist globally: the portfolio must now vary with the base coordinate, and neutrality fails.

The two mechanisms revisited

The fiber bundle picture clarifies why the location-scale and stochastic volatility generalisations (Section 6.3) rest on different structures. The location-scale result is a statement about the *base*: the translation symmetry $x \to x - \tau_w t$ holds regardless of the noise distribution, because it is a property of the tax mechanism (a deterministic multiplicative drain). No reference to the fiber or to preferences is needed.

The stochastic volatility result is a statement about the *connection*: under CRRA, the connection remains flat even when the base is enlarged from x to (x, v). The new coordinate v adds dimensionality to the state space and introduces hedging demand into the portfolio, but the CRRA homogeneity ensures that $\mathbf{w}^*$ remains independent of x. The tax, acting only along x, leaves the (still flat) connection invariant.

Acknowledgements

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

References

Jean-Philippe Bouchaud. The inelastic market hypothesis: A microstructural interpretation. *Quantitative Finance*, 22(10):1785–1795, 2022. doi: 10.1080/14697688.2022.2068052.

Jean-Philippe Bouchaud and J. Doyne Farmer. Self-fulfilling prophecies, quasi non-ergodicity & wealth inequality. SSRN working paper, 2021.

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

George Chacko and Luis M. Viceira. Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets. *The Review of Financial Studies*, 18(4):1369–1402, 2005. doi: 10.1093/rfs/hhi035.

John H. Cochrane. *Asset Pricing*. Princeton University Press, revised edition, 2005.

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

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

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

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 and Ralph S. J. Koijen. In search of the origins of financial fluctuations: The inelastic markets hypothesis. Working paper, Harvard University and University of Chicago, 2022.

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

Steven L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2):327–343, 1993. doi: 10.1093/rfs/6.2.327.

Kirill Ilinski. *Physics of Finance: Gauge Modelling in Non-Equilibrium Pricing*. John Wiley & Sons, 2001.

Harry Kesten. Random difference equations and renewal theory for products of random matrices. *Acta Mathematica*, 131:207–248, 1973.

Roberto Livi and Paolo Politi. *Nonequilibrium Statistical Physics: A Modern Perspective*. Cambridge University Press, 2nd edition, 2017.

Robert C. Merton. An intertemporal capital asset pricing model. *Econometrica*, 41(5):867–887, 1973. doi: 10.2307/1913811.

Mikio Nakahara. *Geometry, Topology and Physics*. CRC Press, 2nd edition, 2003.

Ole Peters. The ergodicity problem in economics. *Nature Physics*, 15:1216–1221, 2019.

Hannes Risken. *The Fokker–Planck Equation: Methods of Solution and Applications*. Springer, 2nd edition, 1989.

Victor M. Yakovenko and J. Barkley Rosser, Jr. Colloquium: Statistical mechanics of money, wealth, and income. *Reviews of Modern Physics*, 81(4):1703–1725, 2009.

Robert Zwanzig. *Nonequilibrium Statistical Mechanics.* Oxford University Press, 2001.