

Extensions to the Wealth Tax Neutrality Framework

Anders G Frøseth*

Draft — 2 March 2026

Abstract

Frøseth (2026) shows that a proportional wealth tax on market values is neutral with respect to portfolio choice, Sharpe ratios, and equilibrium prices under CRRA preferences and geometric Brownian motion. This paper investigates the robustness of that result along two dimensions. First, we extend the neutrality frontier: portfolio neutrality—including all intertemporal hedging demands—is preserved under stochastic volatility (Heston and general Markov diffusions) and Epstein–Zin recursive utility, but breaks under non-homothetic preferences such as HARA. Second, we identify four channels through which implemented wealth taxes depart from neutrality even under CRRA: non-uniform assessment across asset classes, general equilibrium price effects in inelastic markets, progressive threshold structures, and endogenous labour supply. Each channel is formalised and, where possible, calibrated to the Norwegian wealth tax system. The progressive threshold introduces a tax shield that *increases* risk-taking near the exemption boundary—an effect opposite in sign to the HARA distortion—and, at the extreme, generates a participation margin at which investors exit the tax jurisdiction entirely. We formalise this tax-induced migration as the extreme response at the progressive threshold and examine the Norwegian post-2022 experience as a case study. The full framework is applied to evaluate the Saez–Zucman proposal for a global minimum wealth tax on billionaires and the related French proposal for a national minimum tax above €100 million.

JEL Classification: G11, G12, H21, H24, H26, H73.

Keywords: Wealth tax, portfolio choice, tax neutrality, stochastic volatility, HARA preferences, progressive taxation, inelastic markets, tax-induced migration, Saez–Zucman proposal, Norway.

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

Contents

1	Introduction	4
2	Stochastic Volatility	5
2.1	The Heston model	5
2.2	Wealth dynamics with a wealth tax	6
2.3	The Merton problem under Heston	6
2.4	CRRA value function and optimal portfolio	6
2.5	Consumption and welfare	8
2.6	Generalisation to Markov diffusion models	8
2.7	When does neutrality break?	9
2.8	Time-scale dependence of stylised facts	10
3	Beyond CRRA Preferences	10
3.1	HARA utility	10
3.2	Optimal portfolio with a wealth tax	11
3.3	Portfolio distortion	12
3.4	Broader preference specifications	13
3.5	Empirical evidence on wealth-dependent risk-taking	14
4	Non-Uniform Taxation	15
4.1	Institutional setting: the Norwegian wealth tax	16
4.2	Portfolio problem with asset-class-specific assessment	16
4.3	Calibration to the Norwegian system	17
4.4	Leverage and debt deductibility	18
4.5	Empirical evidence	19
4.6	Interaction with non-CRRA preferences	19
4.7	Policy implications	20
5	Inelastic Markets and General Equilibrium	21
5.1	From partial to general equilibrium	21
5.2	The inelastic markets hypothesis	22
5.3	Price impact of wealth-tax-induced flows	23
5.4	Calibration	23
5.5	Implications for neutrality	25
6	Progressive Taxation and Threshold Effects	26
6.1	General framework	26
6.2	Portfolio distortion under CRRA	27
6.3	Norwegian calibration (2026)	28
6.4	Interaction with non-homothetic preferences	29
6.5	Threshold bunching and behavioural responses	29
6.6	Policy implications	30

7	Endogenous Labour Supply and Entrepreneurial Effort	31
7.1	Wealth effects under proportional taxation	31
7.2	The threshold notch and labour supply	32
7.3	Entrepreneurial effort and human capital	32
7.4	Empirical evidence	33
7.5	Policy implications	34
8	Application: A Global Minimum Wealth Tax	34
8.1	The global proposal: assessment through the extension framework	34
8.2	The French variant: a national minimum wealth tax	36
8.3	Comparison with existing wealth taxes	37
8.4	Implementation challenges	38
9	Tax-Induced Migration	38
9.1	Migration as a threshold response	38
9.2	The Norwegian case: 2022–2024	40
9.3	Decomposing behavioural responses	42
9.4	Implications for the framework	42
10	Conclusion	43
A	Leverage and debt deductibility: historical examples	44
B	Time-scale dependence of stylised facts	45
C	Empirical evidence on non-uniform assessment	45

1 Introduction

This paper builds on the wealth tax neutrality framework developed in [Frøseth \(2026\)](#), which established that a proportional wealth tax on market values is neutral with respect to portfolio choice, Sharpe ratios, and equilibrium prices. That analysis proceeds under geometric Brownian motion and then extends to the location-scale family of return distributions. Here we investigate whether the neutrality results survive under more general conditions, and systematically identify the channels through which real-world wealth taxes depart from neutrality.

The paper has three parts. The first (Sections [2–3](#)) extends the neutrality frontier: we show that portfolio neutrality survives under stochastic volatility, general Markov diffusion dynamics, and Epstein–Zin recursive utility, while it breaks under non-homothetic preferences such as HARA utility. The second (Sections [4–7](#)) analyses the channels through which implemented wealth taxes depart from neutrality, even when investors have CRRA preferences: non-uniform assessment across asset classes, general equilibrium price effects in inelastic markets, progressive threshold structures, and endogenous labour supply. The third part applies and extends the framework. Section [8](#) synthesises all channels to evaluate two recent proposals: the Saez–Zucman global minimum wealth tax on billionaires and the French national minimum tax above €100 million. Section [9](#) formalises tax-induced migration as the extreme participation response at the progressive threshold, using the Norwegian post-2022 experience as a case study.

Section [2](#) considers stochastic volatility models, where the variance of returns is itself a random process. Using the Heston model as the leading example, we show that portfolio neutrality—including the intertemporal hedging demand identified by [Merton \(1973\)](#)—is preserved under CRRA preferences. The result generalises to any Markov diffusion model of asset returns. Section [3](#) identifies CRRA as the operative condition: non-homothetic preferences (HARA, wealth-in-utility) break neutrality, while Epstein–Zin preferences preserve it. Section [4](#) derives the portfolio distortion from asset-class-specific assessment discounts, calibrated to the Norwegian system. Section [5](#) analyses general equilibrium price effects through the inelastic markets hypothesis. Section [6](#) formalises the distortion from progressive taxation (thresholds and brackets), and Section [7](#) extends the framework to endogenous labour supply and entrepreneurial effort. Section [8](#) synthesises all channels in a comparative evaluation of the Saez–Zucman global proposal and the French national variant. Section [9](#) formalises tax-induced migration as the extreme participation response at the progressive threshold and examines the Norwegian post-2022 experience as a case study.

Table [1](#) provides an overview of the main results.

Section	Result	Proposition	Sign
<i>Part I: Neutrality frontier</i>			
2	Stochastic volatility (Heston, Markov)	1, 2	Neutral
3	Epstein–Zin recursive utility	4	Neutral
3	HARA (non-homothetic) preferences	3	$\Delta w^* < 0$
<i>Part II: Non-neutrality channels</i>			
4	Non-uniform assessment	5	Asset-dependent
5	Inelastic markets (GE price impact)	—	$\Delta P/P < 0$
6	Progressive threshold (tax shield)	6	$\Delta w^* > 0$
7	Endogenous labour supply	7	Separable [†]
<i>Part III: Applications and extensions</i>			
8	Zucman/French proposals	—	Synthesis
9	Tax-induced migration	—	Exit at W_i^*

[†]Separable under proportional taxation; distorted at progressive thresholds.

Table 1: Summary of main results. “Sign” indicates the direction of the portfolio or price distortion relative to the no-tax benchmark under CRRA preferences.

2 Stochastic Volatility

The neutrality results in Frøseth (2026) are derived first under geometric Brownian motion (constant μ , σ) and then generalised to the location-scale family of return distributions. A natural question is whether the results extend to stochastic volatility models, where the variance of returns is itself a random process. This is empirically important: asset returns exhibit time-varying volatility, fat tails, and leverage effects that are absent under GBM (Heston, 1993; Drăgulescu and Yakovenko, 2002).

We show that, under CRRA preferences, portfolio neutrality extends to the Heston stochastic volatility model—and, more generally, to any Markov diffusion model of asset returns—*without* requiring the location-scale property. The mechanism is different from the location-scale argument: it relies on the homogeneity of the CRRA value function in wealth, which makes optimal portfolio weights (including the intertemporal hedging demand identified by Merton 1973) independent of the wealth level.

2.1 The Heston model

The Heston model specifies the price of a single risky asset and its instantaneous variance as a pair of coupled diffusions:

$$\frac{dS}{S} = \mu dt + \sqrt{v_t} dW_t^{(1)}, \quad (1)$$

$$dv_t = \lambda(\theta - v_t) dt + \kappa\sqrt{v_t} dW_t^{(2)}, \quad (2)$$

where $v_t = \sigma_t^2$ is the instantaneous variance, $\theta > 0$ is the long-run mean variance, $\lambda > 0$ is the rate of mean reversion, $\kappa > 0$ is the volatility of variance, and $\text{corr}(dW^{(1)}, dW^{(2)}) = \rho$. When $\rho < 0$, negative returns coincide with rising volatility—the leverage effect. The Feller condition $2\lambda\theta \geq \kappa^2$ ensures the variance remains strictly positive.

Remark (Heston returns and the location-scale family). [Drăgulescu and Yakovenko \(2002\)](#) derive the probability density of log-returns under the Heston model in closed form and show that it involves a modified Bessel function K_1 of a scaled argument. The distribution exhibits exponential (rather than Gaussian) tails for large returns, and asymmetry when $\rho \neq 0$. These features place Heston returns *outside* the location-scale family. The generalised Propositions 2' and 3' in [Frøseth \(2026\)](#), which require the location-scale property, therefore do not apply directly.

2.2 Wealth dynamics with a wealth tax

An investor allocates a fraction w of wealth to the risky asset and the remainder $1 - w$ to a risk-free asset earning a continuous rate r_f . Following the continuous-time formulation in [Frøseth \(2026\)](#), the proportional wealth tax τ_w enters as an additional drain on the drift of wealth. The investor's wealth evolves as

$$dW = \{W[r_f + w(\mu - r_f) - \tau_w] - C\} dt + wW\sqrt{v_t} dW_t^{(1)}, \quad (3)$$

where $C \geq 0$ is the consumption flow rate. As in the GBM case, the tax reduces the drift of wealth by $\tau_w W$ but leaves the diffusion coefficient—and hence the per-share return dS/S —unchanged. The variance state variable v_t evolves according to (2), independently of the tax.

2.3 The Merton problem under Heston

Consider an investor with CRRA utility over consumption,

$$U(C) = \frac{C^{1-\gamma}}{1-\gamma}, \quad \gamma > 0, \gamma \neq 1, \quad (4)$$

who maximises expected discounted lifetime utility $E\left[\int_0^T e^{-\delta t} U(C_t) dt\right]$ subject to the wealth dynamics (3) and the variance dynamics (2).

The value function is

$$J(W, v, t) = \max_{C, w} E_t \left[\int_t^T e^{-\delta(s-t)} U(C_s) ds \right]. \quad (5)$$

By Bellman's principle, J satisfies the Hamilton–Jacobi–Bellman (HJB) equation

$$0 = \max_{C, w} \left\{ U(C) - \delta J + J_t + J_W [W(r_f + w(\mu - r_f) - \tau_w) - C] \right. \\ \left. + J_v \lambda(\theta - v) + \frac{1}{2} J_{WW} w^2 v W^2 + \frac{1}{2} J_{vv} \kappa^2 v + J_{Wv} w W \kappa v \rho \right\}. \quad (6)$$

2.4 CRRA value function and optimal portfolio

Under CRRA utility, the homogeneity of U suggests a separable value function of the form

$$J(W, v, t) = \frac{W^{1-\gamma}}{1-\gamma} f(v, t), \quad (7)$$

where $f(v, t) > 0$ encodes the dependence on the volatility state and the investment horizon. The partial derivatives are

$$J_W = W^{-\gamma} f, \quad J_{WW} = -\gamma W^{-\gamma-1} f, \quad J_{Wv} = W^{-\gamma} f_v. \quad (8)$$

The first-order condition for the portfolio weight w in (6) is

$$J_W W(\mu - r_f) + J_{WW} w v W^2 + J_{Wv} w W \kappa v \rho = 0. \quad (9)$$

Substituting (8) and dividing through by $W^{1-\gamma} f$:

$$(\mu - r_f) - \gamma w v + \frac{f_v}{f} \kappa v \rho = 0. \quad (10)$$

Solving for w^* :

$$w^* = \underbrace{\frac{\mu - r_f}{\gamma v}}_{\text{myopic demand}} + \underbrace{\frac{f_v}{f} \cdot \frac{\kappa \rho}{\gamma}}_{\text{hedging demand}} \quad (11)$$

Proposition 1 (Portfolio neutrality under stochastic volatility). *Under the Heston model (1)–(2) with CRRA preferences (4) and a proportional wealth tax on all assets, the optimal portfolio weight (11) is independent of the wealth tax rate τ_w .*

Proof. The optimal weight w^* in (11) depends on $\mu, r_f, \gamma, v, \kappa, \rho$, and the ratio f_v/f . It does not depend on the wealth level W . The wealth tax τ_w enters the wealth dynamics (3) and hence the HJB equation (6) only through the drift of W . Under the separable form (7), the function $f(v, t)$ satisfies a PDE obtained by substituting $J = W^{1-\gamma} f/(1-\gamma)$ and the optimal controls into (6). In this PDE, τ_w appears only in the constant (non- v -dependent) term of the drift, alongside r_f . Specifically, the PDE takes the form

$$0 = f_t + \lambda(\theta - v)f_v + \frac{1}{2}\kappa^2 v f_{vv} + h(v; \gamma, \mu, r_f, \kappa, \rho) f + (1 - \gamma)(r_f - \tau_w)f + g(f), \quad (12)$$

where h collects the v -dependent terms arising from the optimal portfolio and $g(f)$ is the contribution from optimal consumption. The wealth tax rate τ_w appears only in the term $(1 - \gamma)(r_f - \tau_w)f$, which shifts the “effective discount rate” but does not interact with v .

Using the standard exponential-affine guess $f(v, t) = \exp(A(t) + B(t)v)$ (Chacko and Viceira, 2005), the Riccati equation for $B(t)$ arises solely from the v -dependent terms in (12) and is therefore independent of τ_w . The ODE for $A(t)$ absorbs the constant terms including τ_w . Consequently, $f_v/f = B(t)$ is independent of τ_w , and the hedging demand $\frac{f_v}{f} \cdot \frac{\kappa \rho}{\gamma}$ is tax-invariant.

Since neither the myopic demand nor the hedging demand depends on τ_w , the full optimal portfolio weight w^* is independent of the wealth tax rate. \square

Remark (Interpretation). The result has an intuitive economic interpretation. Under CRRA, preferences are homogeneous of degree $1 - \gamma$ in wealth. The wealth tax scales the investor’s

wealth by a deterministic factor without altering the return per unit of wealth invested or the dynamics of the volatility state. Since CRRA portfolio weights depend on relative risk-return trade-offs—not on the absolute wealth level—the tax is invisible to the portfolio decision.

This mechanism is distinct from the location-scale argument used in Frøseth (2026) for general utility. There, neutrality follows from a homothetic contraction of the opportunity set that preserves Sharpe ratios. Here, neutrality follows from the homogeneity of preferences, which makes the opportunity set’s *shape* (not its scale) the sole determinant of portfolio choice.

2.5 Consumption and welfare

The first-order condition for consumption in (6) gives

$$C^* = W f(v, t)^{-1/\gamma}. \quad (13)$$

The consumption-wealth ratio $C^*/W = f(v, t)^{-1/\gamma}$ depends on v and t but not on W directly. However, through the function f , the consumption rate depends on τ_w via the effective discount rate in the PDE (12): a higher tax rate reduces the level of f and hence raises the consumption-wealth ratio (the investor consumes a larger fraction of a smaller expected wealth). The tax therefore affects the *level* of consumption and lifetime welfare—only the *portfolio composition* is neutral.

2.6 Generalisation to Markov diffusion models

The argument in Section 2.4 relies on two properties: (i) the wealth tax enters multiplicatively in the wealth dynamics and additively in the drift, and (ii) the CRRA value function is separable in W and the state variables. Neither property is specific to the Heston model.

Consider a general Markov diffusion model with K risky assets and M state variables $\mathbf{X}_t = (X_1, \dots, X_M)^\top$:

$$\frac{dS_i}{S_i} = \mu_i(\mathbf{X}_t) dt + \sum_{j=1}^K \sigma_{ij}(\mathbf{X}_t) dW_t^{(j)}, \quad i = 1, \dots, K, \quad (14)$$

$$dX_m = a_m(\mathbf{X}_t) dt + \sum_{j=1}^M b_{mj}(\mathbf{X}_t) dB_t^{(j)}, \quad m = 1, \dots, M, \quad (15)$$

where the Brownian motions $W^{(j)}$ and $B^{(j)}$ may be correlated. The expected returns μ_i , volatilities σ_{ij} , and state variable dynamics a_m , b_{mj} are all functions of the state \mathbf{X}_t but not of the investor’s wealth.

The investor’s wealth evolves as

$$dW = \{W[r_f(\mathbf{X}_t) + \mathbf{w}^\top(\boldsymbol{\mu}(\mathbf{X}_t) - r_f(\mathbf{X}_t)\mathbf{1}) - \tau_w] - C\} dt + W \mathbf{w}^\top \boldsymbol{\Sigma}(\mathbf{X}_t) d\mathbf{W}_t. \quad (16)$$

Under CRRA utility, the value function takes the form

$$J(W, \mathbf{X}, t) = \frac{W^{1-\gamma}}{1-\gamma} f(\mathbf{X}, t), \quad (17)$$

and the first-order conditions for the optimal portfolio \mathbf{w}^* yield

$$\mathbf{w}^* = \frac{1}{\gamma} \mathbf{V}(\mathbf{X})^{-1} (\boldsymbol{\mu}(\mathbf{X}) - r_f(\mathbf{X})\mathbf{1}) + \frac{1}{\gamma} \mathbf{V}(\mathbf{X})^{-1} \boldsymbol{\Phi}(\mathbf{X}) \frac{\nabla_{\mathbf{X}} f}{f}, \quad (18)$$

where $\mathbf{V}(\mathbf{X}) = \boldsymbol{\Sigma}(\mathbf{X})\boldsymbol{\Sigma}(\mathbf{X})^\top$ and $\boldsymbol{\Phi}(\mathbf{X})$ is a matrix of covariances between asset returns and state variable innovations. Neither term depends on W .

Proposition 2 (Portfolio neutrality under general Markov diffusions). *Let asset returns and state variables follow the Markov diffusion (14)–(15). Under CRRA preferences and a proportional wealth tax on all assets, the optimal portfolio weights (18)—including all intertemporal hedging demands—are independent of the wealth tax rate τ_w .*

This encompasses, as special cases: geometric Brownian motion (constant μ, σ), the Heston stochastic volatility model, the Hull–White model, the SABR model, affine term structure models with stochastic interest rates, and any other model in which returns follow a Markov diffusion with state-dependent coefficients.

Proof. By the same argument as Proposition 1. Under the separable value function (17), the first-order conditions for \mathbf{w}^* involve only the partials J_W , J_{WW} , and J_{WX_m} . The ratios $J_W/(WJ_{WW}) = -1/\gamma$ and $J_{WX_m}/(WJ_{WW}) = -f_{X_m}/(\gamma f)$ are both independent of W . Since τ_w enters the wealth dynamics only through the drift of W , and the PDE for $f(\mathbf{X}, t)$ inherits τ_w only in terms that are independent of \mathbf{X} , the gradient ratio $\nabla_{\mathbf{X}} f/f$ is independent of τ_w . \square

2.7 When does neutrality break?

The stochastic volatility extension identifies CRRA as the operative condition for neutrality beyond the location-scale family. Neutrality can fail when:

1. **Non-CRRA preferences.** If relative risk aversion depends on the wealth level—as with CARA utility, habit formation, or reference-dependent preferences—the portfolio weight becomes a function of W . The wealth tax, by reducing W , then shifts the investor’s risk tolerance and alters the optimal portfolio. Under stochastic volatility, this effect is amplified: the hedging demand depends on the curvature of the value function, which is no longer homogeneous.
2. **Wealth-dependent state dynamics.** If the state variable dynamics depend on the investor’s wealth (e.g., through market impact or endogenous volatility), the tax could alter the state evolution itself. This is empirically unlikely for a single investor but could matter in general equilibrium with a representative agent.
3. **Non-universal taxation.** If the tax applies to some assets but not others, the myopic demand is distorted (as in Frøseth 2026, Section 9). Under stochastic volatility, the hedging

demand may also be distorted if the tax-exempt asset serves as a volatility hedge.

2.8 Time-scale dependence of stylised facts

The stylised facts that motivate the stochastic volatility extension—fat tails, volatility clustering, leverage effects—are not equally pronounced at all horizons. Appendix B reviews the evidence in detail. The key conclusion is that the most dramatic departures from GBM are high-frequency phenomena that attenuate at policy-relevant (monthly to annual) horizons. What *does* persist is time-varying expected returns, regime-like volatility dynamics, and the variance risk premium—precisely the features captured by the general Markov diffusion framework of Section 2.6. The stochastic volatility extension is therefore valuable because it demonstrates that neutrality extends to these empirically relevant features.

3 Beyond CRRA Preferences

The stochastic volatility extension identifies CRRA preferences as the operative condition for portfolio neutrality beyond the location-scale family. Under CRRA, the value function is homogeneous of degree $1 - \gamma$ in wealth, which makes all portfolio demands independent of the wealth level and hence of the tax. This section asks: what happens when preferences depart from CRRA?

The question is empirically motivated. Households with substantial financial wealth often display risk-taking behaviour that is difficult to reconcile with constant relative risk aversion. [Carroll \(2000\)](#) documents that the ultra-wealthy maintain saving rates far exceeding what lifecycle consumption smoothing would predict, suggesting that wealth itself—not just the consumption it finances—enters the objective. [Wachter and Yogo \(2010\)](#) show that household portfolio shares in risky assets *rise* with wealth, a pattern consistent with decreasing relative risk aversion (DRRA) but inconsistent with CRRA. On the other hand, [Brunnermeier and Nagel \(2008\)](#) find that risky portfolio shares respond only weakly to wealth fluctuations, a result more consistent with CRRA or with substantial portfolio inertia.

We take the HARA (Hyperbolic Absolute Risk Aversion) utility class as the tractable leading example. HARA nests CRRA as a special case and generates wealth-dependent portfolio demands through a subsistence consumption parameter. We derive the optimal portfolio under a wealth tax and show that neutrality fails: the tax distorts portfolio composition by raising the present value of subsistence needs.

3.1 HARA utility

The HARA class, analysed by [Merton \(1971\)](#), is defined by the property that the Arrow–Pratt risk tolerance is linear in consumption:

$$T(c) \equiv -\frac{u'(c)}{u''(c)} = \frac{c - \zeta}{\gamma}, \quad (19)$$

where $\zeta \geq 0$ is a subsistence (or habit) level and $\gamma > 0$ is the curvature parameter. The corresponding utility function is

$$u(c) = \frac{(c - \zeta)^{1-\gamma}}{1-\gamma}, \quad c > \zeta, \quad \gamma \neq 1. \quad (20)$$

When $\zeta = 0$, this reduces to CRRA. When $\zeta > 0$, the investor requires a minimum consumption flow ζ per unit time; only consumption in excess of ζ generates the usual isoelastic utility.

The relative risk aversion of u at consumption level c is $\gamma c / (c - \zeta)$, which exceeds γ for all $c > \zeta$ and converges to γ as $c \rightarrow \infty$. Thus HARA with $\zeta > 0$ generates *decreasing relative risk aversion*: wealthier investors, who consume further above subsistence, are effectively less risk-averse.

3.2 Optimal portfolio with a wealth tax

Consider the standard Merton problem under GBM with a single risky asset (drift μ , volatility σ), a risk-free rate r_f , and a proportional wealth tax τ_w . The investor maximises expected discounted lifetime utility subject to the wealth dynamics (3) (with constant σ).

Under HARA utility (20), the value function takes the form

$$J(W, t) = \frac{(W - H)^{1-\gamma}}{1-\gamma} g(t), \quad (21)$$

where H is the *floor wealth*—the present value of the subsistence consumption stream ζ funded entirely by the risk-free asset at the after-tax return $r_f - \tau_w$:

$$H(\tau_w) = \frac{\zeta}{r_f - \tau_w}, \quad r_f > \tau_w. \quad (22)$$

The condition $r_f > \tau_w$ is required for H to be finite: if the wealth tax exceeds the risk-free rate, the investor cannot fund subsistence indefinitely from riskless savings alone.

The investor effectively decomposes wealth into two components: a riskless annuity worth H that funds subsistence, and a *surplus* $S = W - H$ that is invested optimally. The surplus behaves as the wealth of a CRRA investor.

The first-order condition for the portfolio weight in (21) gives the partial derivatives

$$J_W = (W - H)^{-\gamma} g, \quad J_{WW} = -\gamma(W - H)^{-\gamma-1} g. \quad (23)$$

The FOC $J_W W(\mu - r_f) + J_{WW} w \sigma^2 W^2 = 0$ becomes

$$(W - H)(\mu - r_f) - \gamma w \sigma^2 W = 0, \quad (24)$$

yielding the optimal portfolio weight (fraction of total wealth in the risky asset):

$$w^* = \frac{\mu - r_f}{\gamma\sigma^2} \cdot \frac{W - H(\tau_w)}{W} = w_{\text{CRRRA}}^* \cdot \left(1 - \frac{H(\tau_w)}{W}\right) \quad (25)$$

3.3 Portfolio distortion

Proposition 3 (Portfolio distortion under HARA preferences). *Under HARA utility (20) with subsistence level $\zeta > 0$ and a proportional wealth tax τ_w on all assets, the optimal portfolio weight (25) is strictly decreasing in τ_w for all $W > H(\tau_w)$. The tax distortion relative to the zero-tax case is*

$$\Delta w^* \equiv w^*(\tau_w) - w^*(0) = -\frac{\mu - r_f}{\gamma\sigma^2} \cdot \frac{\zeta \tau_w}{W r_f (r_f - \tau_w)} < 0. \quad (26)$$

The magnitude of the distortion is larger for investors with (i) higher subsistence needs ζ , (ii) lower wealth W , (iii) lower risk-free rates r_f , and (iv) higher existing tax rates τ_w . The distortion vanishes as $W \rightarrow \infty$ or $\zeta \rightarrow 0$.

Proof. From (25), $w^* = w_{\text{CRRRA}}^* \cdot (1 - H(\tau_w)/W)$. Since $w_{\text{CRRRA}}^* = (\mu - r_f)/(\gamma\sigma^2)$ is independent of τ_w and W , the tax dependence arises entirely through $H(\tau_w) = \zeta/(r_f - \tau_w)$. Differentiating:

$$\frac{\partial w^*}{\partial \tau_w} = -\frac{w_{\text{CRRRA}}^*}{W} \cdot \frac{\partial H}{\partial \tau_w} = -\frac{w_{\text{CRRRA}}^*}{W} \cdot \frac{\zeta}{(r_f - \tau_w)^2} < 0. \quad (27)$$

The distortion (26) follows from $\Delta w^* = -w_{\text{CRRRA}}^* \cdot (H(\tau_w) - H(0))/W$ with $H(\tau_w) - H(0) = \zeta\tau_w/(r_f(r_f - \tau_w))$. \square

Figure 1 illustrates the mechanism. Under CRRA, the optimal risky share is constant across wealth levels (dashed line). Under HARA, the risky share increases with wealth and approaches the CRRA value from below. The wealth tax amplifies the distortion by raising the floor wealth H , compressing the surplus $W - H$ available for risk-taking. The shaded region shows the additional distortion attributable to the tax. Two features are immediate: the distortion is largest for investors near the subsistence floor and vanishes for the ultra-wealthy; and each curve is defined only for $W > H(\tau_w)$, so higher tax rates shrink the domain of feasible portfolios.

Remark (Effective risk aversion). The result can be restated in terms of an effective relative risk aversion. Under the HARA value function (21), the relative risk aversion with respect to total wealth is

$$\text{RRA}_{\text{eff}}(W) = -\frac{W J_{WW}}{J_W} = \frac{\gamma W}{W - H(\tau_w)} = \frac{\gamma}{1 - H(\tau_w)/W}. \quad (28)$$

This exceeds γ whenever $H > 0$ and is increasing in τ_w . The wealth tax acts as a “fiscal amplifier” of risk aversion: by raising the floor wealth H , it compresses the surplus $W - H$ and pushes the investor toward more conservative portfolios. The amplification is strongest for investors whose wealth is close to the subsistence floor.

Remark (Dollar amount vs. portfolio share). While the portfolio *share* w^* depends on τ_w , the

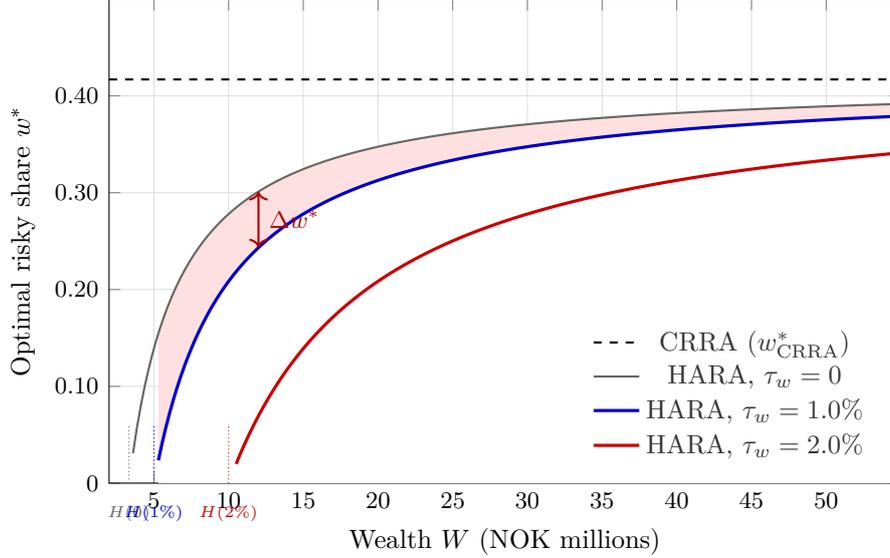


Figure 1: Portfolio distortion under HARA preferences. The dashed line shows the CRRA benchmark (constant across wealth). Solid curves show the optimal risky share under HARA utility for three wealth tax rates. The shaded region between the $\tau_w = 0$ and $\tau_w = 1\%$ curves is the tax-induced distortion Δw^* . Each curve begins at the floor wealth $H(\tau_w) = \zeta/(r_f - \tau_w)$; higher taxes raise H and shrink the feasible domain. Calibration: $\mu - r_f = 5\%$, $\sigma = 20\%$, $\gamma = 3$, $\zeta = 100,000$ NOK, $r_f = 3\%$.

dollar amount invested in the risky asset is $w^*W = w_{\text{CRRA}}^* \cdot (W - H)$ —the CRRA-optimal fraction of the surplus. The tax reduces the risky-asset dollar holding through two channels: it reduces W (direct wealth drain) and increases H (higher subsistence cost), both of which shrink $W - H$.

3.4 Broader preference specifications

The HARA analysis illustrates a general principle: portfolio neutrality requires preferences to be *homothetic in wealth*—that is, the value function must be homogeneous of some degree in W , so that portfolio weights are scale-free. Several empirically motivated preference classes violate this condition.

Wealth in the utility function. Bakshi and Chen (1996) propose a “spirit of capitalism” model in which utility depends on both consumption and wealth: $U = U(C, W)$. If wealth enters because it confers social status, power, or security beyond its consumption value, then the marginal value of wealth is not solely determined by the consumption it can finance. A wealth tax in this framework directly reduces the argument of utility, not merely the budget constraint. The optimal portfolio depends on the cross-partial U_{CW} , which generically makes the portfolio weight a function of W and hence of τ_w . Carroll (2000) provides empirical support: the saving rates of the ultra-wealthy are far too high to be explained by consumption smoothing alone, consistent with wealth entering utility directly.

Epstein–Zin recursive utility. Epstein and Zin (1989) separate the coefficient of relative risk aversion γ from the elasticity of intertemporal substitution ψ , which are constrained to satisfy $\psi = 1/\gamma$ under CRRA.

Proposition 4 (Neutrality under Epstein–Zin). *Under Epstein–Zin preferences with risk aversion γ and EIS $\psi \neq 1/\gamma$, the value function remains homogeneous of degree $1 - \gamma$ in W . The optimal portfolio weights are therefore independent of W , and a proportional wealth tax is portfolio-neutral.*

Proof. The Epstein–Zin aggregator over consumption and continuation value is homogeneous of degree one in (C, V) . Under homothetic budget dynamics (as with a proportional wealth tax that scales all returns uniformly), the value function inherits the homogeneity of the CRRA kernel: $J(W, \mathbf{X}) = W^{1-\gamma} f(\mathbf{X})$. The first-order condition for portfolio weights depends only on \mathbf{X} , not on W . \square

This provides a sharp boundary: departing from CRRA in the direction of Epstein–Zin preserves neutrality (the separation of γ and ψ matters for the consumption response and welfare cost, but not for portfolio composition), while departing in the direction of HARA or wealth-in-utility breaks it.

Non-homothetic consumption. Wachter and Yogo (2010) model households as consuming both necessities and luxuries, with income elasticities that differ across goods. As wealth rises, a smaller fraction is devoted to necessities and the marginal dollar is allocated to luxury consumption, which the household is more willing to put at risk. This generates DRRA and wealth-dependent portfolio shares through a mechanism related to—but distinct from—the HARA subsistence parameter. A wealth tax in the Wachter–Yogo framework would shift the consumption mix toward necessities and reduce the risky portfolio share, qualitatively consistent with the HARA result.

3.5 Empirical evidence on wealth-dependent risk-taking

The magnitude of the portfolio distortion under non-CRRA preferences depends on how strongly risk-taking responds to wealth in practice. The empirical evidence is mixed but informative.

Brunnermeier and Nagel (2008) use panel data on US household portfolios (PSID) and find that the risky asset share responds only weakly to wealth fluctuations driven by income shocks. Their estimates are consistent with CRRA or with DRRA combined with substantial portfolio inertia. If CRRA is approximately correct, the distortions derived above are small.

Calvet and Sodini (2014), using comprehensive Swedish registry data, estimate a financial wealth elasticity of the risky share of approximately 0.2: a 10% increase in financial wealth raises the risky share by about 2 percentage points. This suggests moderate DRRA. Wachter and Yogo (2010) report a similar elasticity when calibrating their non-homothetic model to US data.

These estimates can be used to gauge the HARA distortion. The wealth elasticity of w^* in (25) is

$$\varepsilon \equiv \frac{\partial \ln w^*}{\partial \ln W} = \frac{H(\tau_w)}{W - H(\tau_w)}, \quad (29)$$

which is positive (confirming DRRA) and decreasing in W . Matching $\varepsilon \approx 0.2$ at a representative wealth level gives an estimate of H/W and hence the subsistence-to-wealth ratio that parameterises the distortion. At $\varepsilon = 0.2$, the implied ratio is $H/W \approx 0.17$, meaning roughly one-sixth of wealth is committed to subsistence. The portfolio distortion from a 1% wealth tax at this calibration would be modest but nonzero—reducing w^* by approximately $0.2 \times \tau_w/r_f$ in relative terms.

Implications for policy. The key insight is that the portfolio distortion under non-CRRA preferences is *regressive in wealth*: it is largest for investors near the subsistence floor and vanishes for the ultra-wealthy. This creates a tension with the distributional objectives of wealth taxation. The investors most affected by the portfolio distortion are not the ultra-rich (who are well approximated by CRRA) but those in the lower tail of the wealth distribution subject to the tax—precisely the group for whom the tax is least intended. The ultra-wealthy, whose behaviour motivates much of the policy debate (Carroll, 2000), are paradoxically the group for whom the CRRA neutrality result is most accurate.

4 Non-Uniform Taxation

Sections 2 and 3 examined conditions under which neutrality holds or fails as a property of the return dynamics and the investor’s preferences. We now turn to a different source of non-neutrality: the institutional design of the tax itself. This and the following sections analyse features of real-world wealth taxes—non-uniform assessment, market inelasticity, progressive thresholds, and endogenous effort—that break neutrality even when the investor has CRRA preferences and returns follow a location-scale distribution.

The neutrality results in Frøseth (2026) and in Section 2 above require the wealth tax to apply *uniformly* to all assets: every dollar of wealth is taxed at the same rate regardless of the asset class in which it is held. In practice, no country satisfies this condition. Norway and Switzerland, the two OECD countries with active net wealth taxes, both apply asset-class-specific valuation discounts (Norwegian: *verdsettingsrabatt*) that create differential effective tax rates across asset classes (OECD, 2018; Skatteetaten, 2026).

This section derives the portfolio distortion that arises when the wealth tax is non-uniform. The result is a closed-form expression for the tilt in optimal portfolio weights as a function of the assessment differentials, and it connects directly to the empirical findings of Ring (2024) and Fagereng et al. (2024).

4.1 Institutional setting: the Norwegian wealth tax

Norway levies an annual net wealth tax (*formuesskatt*) at a combined municipal and state rate of 1.0% on net wealth exceeding NOK 1.9 million (2026), rising to 1.1% above NOK 21.5 million. The defining feature of the system is that different asset classes are assessed at different fractions of their market value. As of 2026, the principal assessment fractions are (Skatteetaten, 2026):

Table 2: Norwegian wealth tax assessment fractions (2026)

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

The assessment fractions have changed substantially over time. Before 2023, the discount on shares and commercial property was 45%; it was tightened to 20% for 2023 onwards. The housing assessment model was overhauled in 2010, replacing historical production-cost estimates with hedonic regression models based on market transactions (Ring, 2024). The 75% discount on primary housing was introduced simultaneously to offset the resulting increase in assessed values, a political compromise that embedded a large pro-homeownership tilt into the tax base.

4.2 Portfolio problem with asset-class-specific assessment

Consider K risky assets and one risk-free asset (e.g., bank deposits). The risk-free asset has assessment fraction α_0 and the risky assets have assessment fractions $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_K)^\top$. The statutory wealth tax rate is τ_w .

The after-tax return on the risk-free asset is $r_f - \tau_w \alpha_0$. The after-tax excess return on risky asset i over the risk-free asset is

$$(\mu_i - \tau_w \alpha_i) - (r_f - \tau_w \alpha_0) = (\mu_i - r_f) - \tau_w (\alpha_i - \alpha_0). \quad (30)$$

The term $\tau_w (\alpha_i - \alpha_0)$ is the *tax wedge*: it reduces (increases) the effective excess return of asset i when $\alpha_i > \alpha_0$ ($\alpha_i < \alpha_0$).

The investor's wealth evolves as

$$dW = \{W[r_f - \tau_w \alpha_0 + \mathbf{w}^\top (\boldsymbol{\mu} - r_f \mathbf{1} - \tau_w (\boldsymbol{\alpha} - \alpha_0 \mathbf{1}))] - C\} dt + W \mathbf{w}^\top \boldsymbol{\Sigma} d\mathbf{W}_t. \quad (31)$$

Under CRRA utility and GBM (constant $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$), the first-order conditions for the optimal portfolio yield

$$\boxed{\mathbf{w}^* = \frac{1}{\gamma} \mathbf{V}^{-1} (\boldsymbol{\mu} - r_f \mathbf{1} - \tau_w (\boldsymbol{\alpha} - \alpha_0 \mathbf{1}))} \quad (32)$$

where $\mathbf{V} = \Sigma \Sigma^\top$ is the return covariance matrix.

Proposition 5 (Portfolio distortion under non-uniform taxation). *Under CRRA preferences and a proportional wealth tax with asset-class-specific assessment fractions $\alpha_0, \alpha_1, \dots, \alpha_K$, the optimal portfolio weights (32) depend on the tax rate τ_w unless all assessment fractions are equal. The distortion relative to the uniform-tax case is*

$$\Delta \mathbf{w}^* \equiv \mathbf{w}^*(\boldsymbol{\alpha}) - \mathbf{w}^*(\alpha_0 \mathbf{1}) = -\frac{\tau_w}{\gamma} \mathbf{V}^{-1}(\boldsymbol{\alpha} - \alpha_0 \mathbf{1}). \quad (33)$$

The investor overweights assets with $\alpha_i < \alpha_0$ (those receiving a larger valuation discount than the risk-free asset) and underweights assets with $\alpha_i > \alpha_0$.

Proof. Under uniform assessment $\alpha_i = \alpha_0$ for all i , the tax enters (31) as $-\tau_w \alpha_0 W$, a constant drain on the drift that does not interact with \mathbf{w} . The FOC gives $\mathbf{w}_{\text{uniform}}^* = (1/\gamma) \mathbf{V}^{-1}(\boldsymbol{\mu} - r_f \mathbf{1})$, independent of τ_w . Under non-uniform assessment, (32) contains the additional term $-(\tau_w/\gamma) \mathbf{V}^{-1}(\boldsymbol{\alpha} - \alpha_0 \mathbf{1})$. Subtracting yields (33). \square

Remark (Sharpe ratio distortion). The non-uniform tax distorts after-tax Sharpe ratios. For a single risky asset with assessment fraction α_1 , the after-tax Sharpe ratio is

$$\text{SR}^{\text{after}} = \frac{\mu - r_f - \tau_w(\alpha_1 - \alpha_0)}{\sigma}. \quad (34)$$

When $\alpha_1 < \alpha_0$ (as for Norwegian equities relative to bank deposits), the after-tax Sharpe ratio *exceeds* the pre-tax ratio. The valuation discount artificially inflates the risk-adjusted return of the tax-advantaged asset, distorting the price signal that the investor faces.

4.3 Calibration to the Norwegian system

The Norwegian assessment fractions in Table 2 can be used to quantify the portfolio distortion. Consider a simplified two-asset economy: listed equities ($\alpha_1 = 0.80$) and bank deposits ($\alpha_0 = 1.00$). At a statutory rate $\tau_w = 1.0\%$, the tax wedge is

$$\tau_w(\alpha_1 - \alpha_0) = 0.01 \times (0.80 - 1.00) = -0.002, \quad (35)$$

i.e., the effective tax on equities is 20 basis points per year lower than on deposits. This is equivalent to a permanent annual subsidy of 0.2% on the equity excess return.

For a CRRA investor with $\gamma = 4$ and equity volatility $\sigma = 0.20$, the portfolio tilt from (33) is

$$\Delta w^* = -\frac{0.01}{4} \cdot \frac{0.80 - 1.00}{0.04} = +\frac{0.002}{0.16} = +1.25\%. \quad (36)$$

The investor increases the equity weight by 1.25 percentage points relative to the neutral benchmark—a modest but non-trivial distortion that compounds over the portfolio and across asset classes.

For primary housing ($\alpha = 0.25$), the tilt is far larger. If housing is modelled as an investable

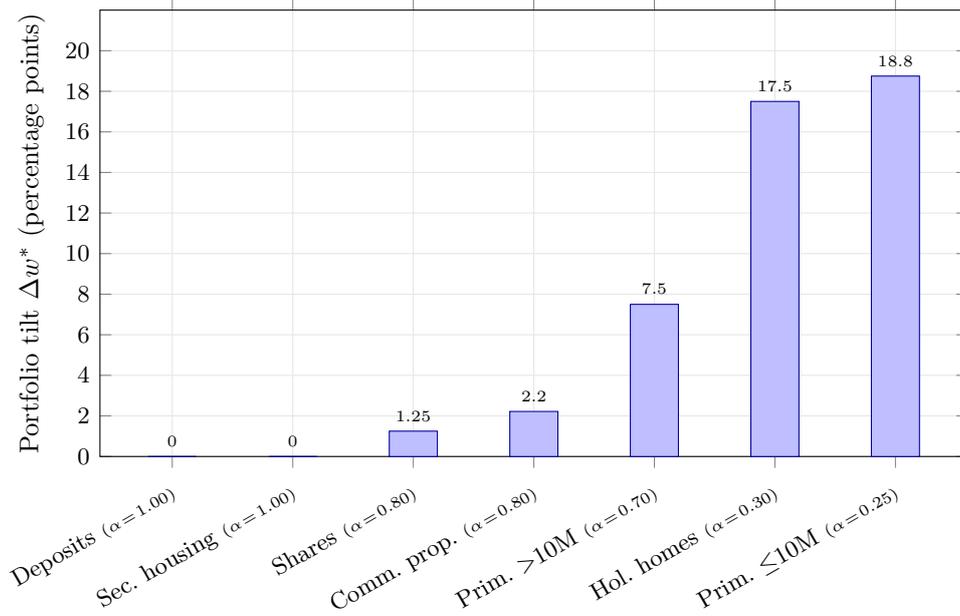


Figure 2: Portfolio tilt Δw^* toward each asset class relative to bank deposits, under the Norwegian assessment system. The tilt is computed from (33) with $\tau_w = 1.0\%$, $\gamma = 4$, and asset-class volatilities $\sigma = 0.20$ (listed shares), $\sigma = 0.15$ (commercial property), $\sigma = 0.10$ (housing classes). Assessment fractions α_i are shown below each bar. The two-order-of-magnitude spread illustrates the strong incentive to hold primary housing over financial assets.

asset with volatility $\sigma_h = 0.10$, the distortion becomes

$$\Delta w_h^* = -\frac{0.01}{4} \cdot \frac{0.25 - 1.00}{0.01} = +18.75\%. \quad (37)$$

This large number reflects both the 75% discount and the low volatility of housing. While the simplified calculation ignores important features of housing (illiquidity, leverage, consumption services), it illustrates the order of magnitude: the Norwegian assessment system creates a powerful incentive to hold primary housing relative to financial assets.

Figure 2 extends the two-asset calibration to the main Norwegian asset classes, using the assessment fractions from Table 2 and asset-class-specific volatility estimates. The resulting portfolio tilts span two orders of magnitude: from zero for fully assessed assets (bank deposits, secondary housing) to nearly 19 percentage points for primary housing below 10M NOK. The large spread reflects the combination of generous valuation discounts and low asset-class volatility in the housing sector.

4.4 Leverage and debt deductibility

The portfolio distortions derived above are amplified by the interaction between valuation discounts and debt deductibility. In the Norwegian system, debt has historically been deductible at face value against the wealth tax base, while assets receiving valuation discounts are assessed below market value. For a leveraged position in asset i with loan-to-value ratio $\ell_i = D_i/V_i$, the

net contribution to taxable wealth per unit of asset value is

$$\frac{\text{Taxable wealth}}{V_i} = \alpha_i - \beta_i \ell_i, \quad (38)$$

where $\beta_i \leq 1$ is the fraction of associated debt that is deductible.

Leverage and effective tax rates. Under full debt deduction ($\beta_i = 1$), the net taxable contribution $\alpha_i - \ell_i$ becomes negative whenever leverage exceeds the assessment fraction. This created a powerful sheltering strategy: highly leveraged positions in discounted asset classes could offset the tax base from other assets (see Appendix A for historical examples and calibrations). Since total net taxable wealth is floored at zero, the mechanism operates by sheltering positive wealth in fully assessed assets.

Norway’s proportional debt reduction (*gjeldsreduksjon*) sets $\beta_i = \alpha_i$ for discounted assets, giving

$$\frac{\text{Taxable wealth}}{V_i} = \alpha_i(1 - \ell_i), \quad (39)$$

which is non-negative for all $\ell_i \leq 1$, eliminating the negative tax base. However, differential effective rates persist. Under full debt deduction, the effective wealth tax rate on a leveraged position with equity $E_i = V_i(1 - \ell_i)$ is

$$\tau_{w,i}^{\text{eff}} = \tau_w \cdot \frac{\alpha_i - \ell_i}{1 - \ell_i}, \quad (40)$$

which can be negative—a qualitatively distinct distortion that creates incentives to leverage into discounted asset classes. The proportional debt reduction rule attenuates this to $\tau_{w,i}^{\text{eff}} = \tau_w \alpha_i$, which remains below τ_w for discounted assets.

4.5 Empirical evidence

The theoretical distortions derived above have empirical counterparts. Appendix C reviews the evidence from Norway (Ring, 2024; Fagereng et al., 2024) and Switzerland (Brühlhart et al., 2022) in detail. The key findings are: (i) portfolio composition is unaffected when the tax does not discriminate between asset classes, consistent with our neutrality result under uniform assessment; (ii) when assessment differentials create distinct after-tax returns, households do rebalance, but slowly—full reoptimisation takes approximately five years; and (iii) reported wealth responses to tax rate changes are large but dominated by mobility and valuation responses rather than real portfolio reallocation. The sluggish adjustment implies that the theoretical distortions derived in this paper represent long-run equilibrium magnitudes; short-run portfolio responses to a newly introduced or reformed wealth tax will be smaller.

4.6 Interaction with non-CRRA preferences

The distortions from non-uniform taxation and from non-CRRA preferences identified in Section 3 can compound. Under HARA utility with subsistence level $\zeta > 0$ and asset-class-specific

assessment, the optimal weight of the single risky asset becomes

$$w^* = \frac{\mu - r_f - \tau_w(\alpha_1 - \alpha_0)}{\gamma\sigma^2} \cdot \frac{W - H(\tau_w, \alpha_0)}{W}, \quad (41)$$

where $H(\tau_w, \alpha_0) = \zeta/(r_f - \tau_w\alpha_0)$ is the floor wealth evaluated at the after-tax risk-free rate. The first factor captures the Sharpe-ratio distortion from non-uniform assessment (Proposition 5); the second captures the surplus-wealth effect from non-CRRA preferences (Proposition 3). The two distortions are multiplicative and can reinforce each other: an investor near the subsistence floor holding an asset with a favourable assessment discount experiences both an inflated effective Sharpe ratio and a compressed risk-taking capacity.

Remark (Stochastic volatility and non-uniform assessment). When return volatility varies across asset classes—equities being substantially more volatile than real estate at monthly frequencies—the hedging demand from the stochastic volatility extension (Section 2) interacts with non-uniform assessment. An investor who hedges against volatility shocks by tilting toward low-volatility assets may find this tilt reinforced or opposed by the tax-induced tilt from assessment differentials. Formalising this interaction requires a multi-asset stochastic volatility model with asset-class-specific dynamics, which we leave for future work.

4.7 Policy implications

The analysis highlights a fundamental tension in wealth tax design. Uniform assessment across all asset classes—the condition required for portfolio neutrality—conflicts with several practical objectives.

First, valuation accuracy varies by asset class. Bank deposits, listed equities, and government bonds have observable market prices; private business equity, real estate, and collectibles do not (OECD, 2018). Assessment discounts are often justified as compensation for valuation uncertainty, but they introduce portfolio distortions as a side effect.

Second, assessment discounts serve political objectives. Norway’s 75% discount on primary housing reflects a policy choice to protect homeownership, not a valuation concern—the hedonic regression model introduced in 2010 provides accurate market values (Ring, 2024). The discount was introduced precisely to offset the increase in assessed values that the improved model would generate.

Third, differential assessment creates an implicit industrial policy: it channels capital toward tax-advantaged asset classes and away from those assessed at full market value. The Norwegian system, which taxes bank deposits at full value while discounting housing by 75%, provides a large incentive to hold housing wealth—an allocation that may reduce productive capital formation.

Fourth, as shown in Section 4.4, valuation discounts interact with debt deductibility to create a leverage channel that amplifies the portfolio distortion. Under full debt deduction, leveraged positions in discounted asset classes can generate negative taxable contributions, allowing investors to shelter wealth held in other asset classes. This incentivises not only a tilt toward discounted

assets but also excessive leverage in those positions—a distortion with potential systemic implications for financial stability. Norway’s proportional debt reduction rules (*gjeldsreduksjon*) eliminate the most egregious arbitrage (negative tax bases), but the remaining differential in effective rates continues to favour leveraged investment in discounted asset classes over unleveraged holdings in fully assessed ones.

From the perspective of neutrality, the first-best policy is uniform assessment at market value ($\alpha_i = 1$ for all i) combined with full debt deductibility at face value. [Scheuer and Slemrod \(2021\)](#) survey the design challenges and note that uniform assessment is feasible for liquid assets but faces significant obstacles for illiquid ones. A second-best alternative is to equalise assessment fractions across liquid financial assets (equities, bonds, deposits) while accepting that illiquid assets require separate treatment. This would preserve portfolio neutrality within the financial portfolio while acknowledging that housing and private business equity present distinct valuation challenges. Any remaining valuation discounts should be paired with proportional debt reduction to prevent the leverage arbitrage documented in [Section 4.4](#).

These portfolio distortions are derived in partial equilibrium: they take asset prices as given. When all wealth-tax payers simultaneously adjust their portfolios, the resulting aggregate flows feed back into equilibrium prices. [Section 5](#) examines the magnitude of this general equilibrium response.

5 Inelastic Markets and General Equilibrium

The preceding sections analyse the wealth tax from the perspective of an individual investor who takes asset returns as given. In this partial equilibrium setting, the CRRA neutrality result states that portfolio *weights* are unaffected by the tax. But the tax reduces the investor’s total wealth, and hence the *dollar amount* allocated to each asset. When all taxed investors simultaneously reduce their dollar holdings, aggregate demand for risky assets falls. This section asks: how do equilibrium prices respond, and does the magnitude of the response depend on the structure of market demand?

The answer, we argue, depends critically on the price elasticity of aggregate equity demand. Under the traditional assumption of elastic markets—where many marginal investors stand ready to absorb demand shifts at close to fundamental value—the price adjustment is modest and efficient. Under the *inelastic markets hypothesis* ([Gabaix and Koijen, 2021](#)), aggregate equity demand is far less elastic than commonly assumed, and even modest flows can produce amplified price effects.

5.1 From partial to general equilibrium

Under the CRRA neutrality result ([Section 2](#), [Section 2.6](#)), the optimal portfolio weight w^* is independent of τ_w . However, the dollar amount invested in equities is w^*W , and W is reduced by the tax. In the long run, an investor facing a perpetual tax τ_w accumulates less wealth by a factor that depends on τ_w relative to the growth rate of wealth.

If a fraction ϕ of total equity market capitalisation is held by taxed investors, the aggregate demand shift induced by the tax is

$$\Delta F = -\tau_w \cdot \phi \cdot P_{\text{eq}}, \quad (42)$$

where P_{eq} is the equilibrium market capitalisation and ΔF represents the annual net flow out of equities required to fund tax payments. We treat ΔF as a flow (per period) rather than a level shift, since the tax recurs annually.

In a frictionless general equilibrium with fully elastic demand, this flow is absorbed efficiently: prices adjust by exactly the amount needed to reflect the lower after-tax wealth, expected returns rise commensurately, and the resulting equilibrium is Pareto-efficient conditional on the tax. The “price impact” is simply the new fundamental value.

5.2 The inelastic markets hypothesis

[Gabaix and Koijen \(2021\)](#) challenge the elastic-demand assumption. Using granular instrumental variables constructed from mutual fund flows, they estimate the price elasticity of aggregate US equity demand and find a striking result: *investing \$1 in the stock market increases the aggregate market’s value by approximately \$5*. The multiplier M , defined as the ratio of the price change to the flow that caused it, is estimated at approximately $M \approx 5$, with a range of 3–8 across specifications.

The low elasticity arises because the marginal holders of equity—index funds, pension funds, insurance companies—operate under mandates that fix their equity allocations within narrow bands. When aggregate demand shifts, few participants can absorb the change, and prices must move substantially to clear the market. [Koijen and Yogo \(2019\)](#) develop the demand-system framework that underpins these estimates, showing that institutional demand is far more price-inelastic than the efficient markets paradigm assumes.

The micro-level foundations are provided by [Bouchaud et al. \(2009\)](#) and [Bouchaud \(2022\)](#). Large institutional orders are split into many small child orders and executed over days to months, creating persistent order flow that moves prices incrementally. Revealed liquidity—the depth of the visible order book near the current price—is extremely low relative to typical order sizes. [Bouchaud \(2022\)](#) derives the aggregate multiplier from a latent order book model: because latent liquidity is sparse near the prevailing price, even moderate flows push through the available depth and generate disproportionate price changes. The empirical signature is the well-documented square-root law of price impact, $I(Q) \propto \sqrt{Q}$, which holds across stocks, futures, and options markets ([Bouchaud et al., 2009](#)).

The fire sales literature provides further micro-level evidence. [Coval and Stafford \(2007\)](#) show that mutual funds experiencing large outflows create measurable price pressure in their commonly held securities, with prices subsequently reverting—consistent with temporary demand-driven dislocations rather than information-driven price changes.

5.3 Price impact of wealth-tax-induced flows

Combining the neutrality framework with the inelastic markets hypothesis yields a distinction between two effects of the wealth tax on asset prices:

Fundamental valuation effect. The tax reduces the after-tax wealth of taxed investors, lowering their lifetime consumption and hence the fundamental value of claims on their future income. In a representative-agent economy, this would reduce equilibrium prices by a factor proportional to τ_w relative to the discount rate. This effect is present in any general equilibrium model and does not depend on market elasticity.

Flow-amplification effect. Under inelastic markets, the aggregate flow ΔF required to fund tax payments generates a price response that *exceeds* the fundamental valuation change. If $\Delta F/P_{\text{eq}}$ is the flow as a fraction of market capitalisation, the price impact is

$$\frac{\Delta P}{P} \approx -M \cdot \frac{\Delta F}{P_{\text{eq}}} = -M \cdot \tau_w \cdot \phi, \quad (43)$$

where M is the Gabaix–Kojien multiplier and ϕ is the share of market capitalisation held by taxed investors. The excess price impact $(M - 1) \cdot \tau_w \cdot \phi$ represents a flow-driven dislocation that does not arise in frictionless markets.

Figure 3 illustrates the mechanism. The supply of equity is fixed (vertical). A wealth tax shifts demand to the left by the flow ΔF . Under elastic demand, the equilibrium price falls modestly from P_0 to P_1^E ; under inelastic demand, the same flow produces a much larger decline to P_1^I . The shaded region represents the excess price impact—the amplification due to market inelasticity.

Remark (Incidence on non-taxed investors). The price decline falls on *all* equity holders, not only those subject to the wealth tax. In a small open economy such as Norway, where foreign investors hold a large share of the equity market, a substantial fraction of the price impact is borne by non-taxed investors. This creates an incidence shift: the effective burden of the wealth tax is partly “exported” through the price channel. Conversely, in a closed economy or one where only domestic investors hold equities, the price decline feeds back into the tax base, reducing future tax revenue.

5.4 Calibration

A back-of-the-envelope calculation illustrates the potential magnitude. Consider a wealth tax at rate $\tau_w = 1\%$ applied to equity holdings that constitute $\phi = 20\%$ of total market capitalisation (the remainder being held by tax-exempt institutions and foreign investors). Under the conservative assumption that the entire tax liability is met by selling equities (refined below), the tax-induced flow as a fraction of market capitalisation is

$$\frac{\Delta F}{P_{\text{eq}}} = \tau_w \cdot \phi = 0.01 \times 0.20 = 0.002 = 0.2\%. \quad (44)$$

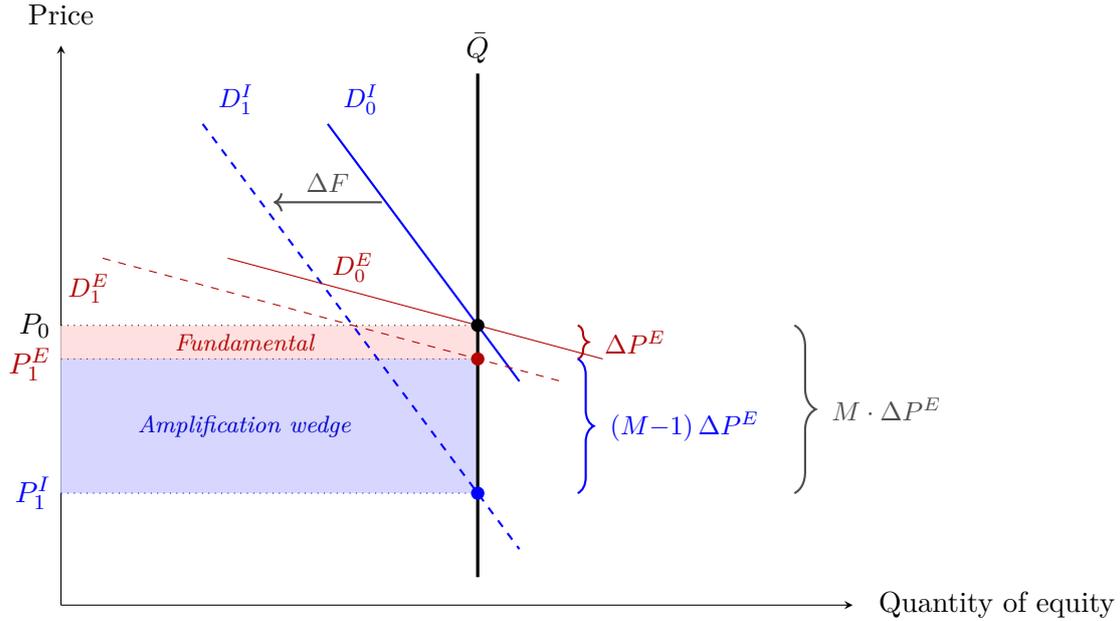


Figure 3: Decomposition of the price impact of a wealth-tax-induced demand shift. Supply is fixed at \bar{Q} . The wealth tax shifts demand left by $\Delta F = \tau_w \phi P_0$. The total price decline from P_0 to P_1^I decomposes into two components: (i) the *fundamental effect* (light red), equal to $\Delta P^E = \tau_w \phi$, which is the decline that would occur under perfectly elastic demand; and (ii) the *amplification wedge* (light blue), equal to $(M-1) \Delta P^E$, which is the excess decline due to market inelasticity. The total impact is M times the fundamental effect, where $M \approx 5$ is the Gabaix–Kojien multiplier. The diagram shows the full tax liability as the demand shift; actual equity outflows are smaller once dividend and liquid-income channels are accounted for (see Section 5.4).

Under frictionless markets ($M = 1$), prices decline by 0.2%—a negligible effect. Under the inelastic markets hypothesis ($M = 5$), prices decline by

$$\frac{\Delta P}{P} = -5 \times 0.002 = -1.0\%. \quad (45)$$

A one percent annual price decline, compounded over the holding period, represents a meaningful reduction in wealth for all equity holders.

Table 3 shows the sensitivity of this estimate to the multiplier across the range reported in the literature.

	Multiplier M		
	3 (low)	5 (central)	8 (high)
$\Delta P/P$ (Norwegian, $\tau_w = 1\%$, $\phi = 0.20$)	-0.6%	-1.0%	-1.6%
$\Delta P/P$ (Zucman, $\tau_w = 2\%$, $\phi_B = 0.08$)	-0.5%	-0.8%	-1.3%

Table 3: Sensitivity of annual equilibrium price impact to the Gabaix–Kojien multiplier M . Central estimate $M = 5$; range 3–8 from Gabaix and Kojien (2021).

Several caveats apply. First, the Gabaix–Kojien multiplier is estimated for the US equity market; no published estimates exist for global equity markets or specific European exchanges. For the

Saez–Zucman global proposal, the relevant multiplier applies to the world equity market. On one hand, greater aggregate depth may lower M relative to the US estimate. On the other, a coordinated global tax eliminates the substitution channel through which selling pressure in one market is absorbed by capital inflows from untaxed jurisdictions, which could sustain or increase M . [Kojen and Yogo \(2019\)](#) estimate a global equity demand elasticity of approximately 1.2 in their demand-system framework, broadly consistent with the US multiplier range, but the mapping to a flow multiplier is not direct. The net effect is ambiguous, and the sensitivity range in [Table 3](#) should be interpreted accordingly. Second, and most importantly, the model assumes that taxed investors liquidate equities to meet their obligations. In practice, investors typically cover the wealth tax through a hierarchy of payment sources: other liquid income (salary, interest, rental income), dividend extractions from their companies, and—only as a last resort—outright share sales (see [Frøseth, 2026](#), for a detailed discussion). [Berzins et al. \(2022\)](#) find that Norwegian private firms increase dividend payouts in response to their owners’ wealth tax obligations, confirming that dividend extraction is the dominant margin. [Bjørneby et al. \(2023\)](#) document the same pattern for closely held firms, where dividend payments rise roughly in proportion to the tax liability. Using Colombian data, [Peydró et al. \(2025\)](#) show that the firm-level response depends on the financial environment: firms whose owners face the wealth tax increase leverage by substituting bank debt for extracted equity, and those with credit access actually increase investment—suggesting that the sign of the real effect is ambiguous and depends on the availability of alternative financing. The effective equity outflow ΔF is therefore substantially smaller than the full tax liability $\tau_w \phi P_0 Q$, and should be interpreted as the residual selling pressure after dividends and liquid assets are exhausted.

Furthermore, in Norway the personal wealth tax is collected through quarterly advance payments (*forskuddsskatt*) during the assessment year, not as a single year-end lump sum. Whatever equity selling does occur is therefore spread over four quarters, reducing the instantaneous flow impact and giving markets time to absorb the pressure gradually. Third, a pre-announced, recurring wealth tax allows markets to anticipate the flow, potentially front-loading the price adjustment. Fourth, the square-root impact law $I(Q) \propto \sqrt{Q}$ suggests that large aggregate flows may have less-than-proportional impact per unit, moderating the linear approximation in [\(43\)](#).

Norway introduced a deferral scheme (*utsettelsesordning*) for wealth tax payments starting in 2026, allowing taxpayers to defer payment for up to three years at a market-rate interest charge. This smooths the liquidity demand and may reduce the flow-amplification effect, though it does not eliminate the fundamental valuation effect.

5.5 Implications for neutrality

The inelastic markets framework introduces a distinction between two levels of neutrality that the partial equilibrium analysis conflates:

1. **Portfolio-weight neutrality.** Under CRRA preferences and uniform assessment, each investor’s optimal portfolio *weights* are independent of τ_w . This result survives in general equilibrium: even after prices adjust, the optimal weight remains $w^* = (\mu' - r_f)/(\gamma\sigma'^2)$ evaluated at the new equilibrium μ', σ' , and the investor’s portfolio share is unchanged.

2. **Price neutrality.** The tax is *not* neutral with respect to equilibrium asset prices. Under inelastic markets, the price impact is amplified by the multiplier M , creating a wedge between the pre-tax and post-tax equilibrium that exceeds the fundamental valuation change.

The distinction matters for policy evaluation. Portfolio-weight neutrality means the tax does not distort relative asset allocation—an important efficiency property. But the lack of price neutrality means the tax can depress asset values, reduce market capitalisation, and shift incidence to non-taxed investors. These general equilibrium effects operate through a channel entirely absent from the partial equilibrium analysis that dominates the wealth tax literature.

The interaction with the extensions developed earlier in this paper creates additional channels. Under non-CRRA preferences (Section 3), the price decline reduces investors’ surplus wealth $W - H$, increasing effective risk aversion and potentially triggering further selling—a feedback loop between the flow-amplification effect and the HARA distortion. Under non-uniform assessment (Section 4), asset-class-specific flows generate differential price impacts across markets, with larger effects in markets for illiquid or concentrated-ownership assets where the effective multiplier may substantially exceed the aggregate estimate.

A full general equilibrium analysis integrating the Merton portfolio framework with the Gabaix–Kojien demand system is beyond the scope of this paper but represents an important direction for future work.

6 Progressive Taxation and Threshold Effects

The neutrality result of Frøseth (2026) rests on the tax being *proportional*: the same rate τ_w applied to every unit of wealth. Every wealth tax implemented in practice departs from proportionality by imposing a positive threshold (the Norwegian *bunnfradrag*) below which no tax is due, and often multiple brackets above it. This section formalises the portfolio distortion created by progressive taxation and calibrates it to the Norwegian system.

6.1 General framework

Consider a progressive wealth tax with K brackets. Let $\bar{W}_1 < \bar{W}_2 < \dots < \bar{W}_K$ be the bracket thresholds and $\tau_1 < \tau_2 < \dots < \tau_K$ the associated marginal rates, with $\tau_0 = 0$ below \bar{W}_1 . The tax liability for an investor with wealth W in bracket j (i.e. $\bar{W}_j < W \leq \bar{W}_{j+1}$, with $\bar{W}_{K+1} = \infty$) is

$$T(W) = \sum_{k=1}^j (\tau_k - \tau_{k-1}) \max(0, W - \bar{W}_k) = \tau_j W - R_j, \quad (46)$$

where the *cumulative rebate*

$$R_j = \sum_{k=1}^j (\tau_k - \tau_{k-1}) \bar{W}_k \quad (47)$$

is the tax saving, relative to a proportional tax at the marginal rate τ_j , generated by the exemption and lower-rate brackets. The decomposition in (46) is the key observation: a progressive

tax equals a proportional tax at the marginal rate minus a lump-sum rebate that depends only on the bracket structure, not on the portfolio.

The effective average rate is

$$\bar{\tau}(W) = \frac{T(W)}{W} = \tau_j - \frac{R_j}{W}, \quad (48)$$

which is strictly increasing in W and converges to the marginal rate τ_j as $W \rightarrow \infty$.

6.2 Portfolio distortion under CRRA

For an investor in bracket j with $W \gg \bar{W}_j$, the wealth dynamics are

$$dW = [(r_f - \tau_j)W + wW(\mu - r_f) + R_j - c]dt + wW\sigma dZ. \quad (49)$$

The proportional component $\tau_j W$ reduces all returns by τ_j , preserving neutrality of the excess return $\mu - r_f$. The rebate R_j enters as a constant income flow, independent of the portfolio. In the Merton continuous-time framework, a constant income flow y is equivalent to a riskless asset with present value $y/(r_f^a)$, where $r_f^a = r_f - \tau_j$ is the after-tax risk-free rate (Merton, 1971).

Definition 1 (Tax shield). The *tax shield* of a progressive wealth tax for an investor in bracket j is the present value of the cumulative rebate:

$$H_\tau = \frac{R_j}{r_f - \tau_j}. \quad (50)$$

Under CRRA preferences, the optimal risky share of financial wealth is determined by the ratio of total wealth (financial plus implicit) to financial wealth alone.

Proposition 6 (Progressive tax distortion). *Under a progressive wealth tax with bracket structure $\{(\bar{W}_k, \tau_k)\}_{k=1}^K$, an investor with CRRA preferences, risk aversion γ , and financial wealth W in bracket j holds a risky portfolio share*

$$w^* = \frac{\mu - r_f}{\gamma\sigma^2} \frac{W + H_\tau}{W} = w_{\text{neutral}}^* \left(1 + \frac{H_\tau}{W}\right), \quad (51)$$

where $w_{\text{neutral}}^* = (\mu - r_f)/(\gamma\sigma^2)$ is the proportional-tax optimal share and H_τ is the tax shield from (50).

Proof. With the constant income flow R_j , the investor's total investable wealth is $W + H_\tau$, of which H_τ is implicitly riskless (the rebate is earned regardless of portfolio returns, conditional on $W > \bar{W}_j$). The CRRA optimal risky share of total wealth is w_{neutral}^* ; expressing this as a share of financial wealth W gives (51). \square

Proposition 6 establishes three results. First, the progressive structure *increases* the risky share relative to the proportional benchmark. The tax shield H_τ acts like a riskless endowment—the exempted wealth generates a certain tax saving each period—which tilts the portfolio toward

risk. Second, the distortion is *decreasing* in wealth:

$$\frac{\partial}{\partial W} \frac{H_\tau}{W} = -\frac{H_\tau}{W^2} < 0,$$

so investors just above the threshold are most affected, and the effect vanishes for $W \rightarrow \infty$. Third, the distortion is *progressive*: it is largest for moderate-wealth investors (who benefit most from the exemption relative to their wealth) and negligible for the very wealthy.

Remark (Direction of distortion). The threshold distortion runs in the *opposite* direction from the HARA distortion of Section 3. There, subsistence consumption creates a floor wealth H that *reduces* risk-taking; here, the tax exemption creates a shield H_τ that *increases* it. Whether the net effect is conservative or aggressive depends on the relative magnitude of H and H_τ .

6.3 Norwegian calibration (2026)

The Norwegian wealth tax for 2026 has two brackets above a threshold (Skatteetaten, 2026):

Bracket	Threshold	Municipal	State	Total rate
Below threshold	$\bar{W}_1 = 1,900,000$		0%	
Bracket 1	\bar{W}_1 to $\bar{W}_2 = 21,500,000$	0.35%	0.65%	$\tau_1 = 1.00\%$
Bracket 2	Above \bar{W}_2	0.35%	0.75%	$\tau_2 = 1.10\%$

Thresholds are per person; married couples receive double.

Bracket 1 investors ($1.9\text{M} < W \leq 21.5\text{M}$ NOK). The cumulative rebate is $R_1 = \tau_1 \cdot \bar{W}_1 = 0.01 \times 1,900,000 = 19,000$ NOK per year. With $r_f = 3\%$:

$$H_\tau = \frac{19,000}{0.03 - 0.01} = 950,000 \text{ NOK}.$$

For an investor with $W = 5\text{M}$ NOK, $H_\tau/W = 19\%$: the progressive structure increases the optimal risky share by 19% relative to the proportional benchmark. At $W = 15\text{M}$ NOK, $H_\tau/W = 6.3\%$.

Bracket 2 investors ($W > 21.5\text{M}$ NOK). The cumulative rebate is $R_2 = R_1 + (\tau_2 - \tau_1) \bar{W}_2 = 19,000 + 0.001 \times 21,500,000 = 40,500$ NOK per year, and

$$H_\tau = \frac{40,500}{0.03 - 0.011} = 2,131,579 \text{ NOK} \approx 2.1\text{M} \text{ NOK}.$$

For $W = 30\text{M}$ NOK, $H_\tau/W = 7.1\%$; for $W = 100\text{M}$, $H_\tau/W = 2.1\%$; for $W = 1\text{B}$, $H_\tau/W = 0.2\%$. The progressive distortion is thus economically significant for moderate wealth but vanishes rapidly at the top of the distribution.

6.4 Interaction with non-homothetic preferences

Combining the progressive tax (Proposition 6) with HARA preferences (Proposition 3 in Section 3) gives a portfolio share that reflects both the subsistence floor and the tax shield:

$$w^* = w_{\text{neutral}}^* \frac{W - H + H_\tau}{W}, \quad (52)$$

where $H = \zeta/(r_f - \tau_j)$ is the HARA floor wealth evaluated at the after-tax risk-free rate of the investor’s bracket, and H_τ is the tax shield from (50).

The two distortions partially offset: subsistence needs push the portfolio toward safety ($-H$), while the tax exemption pushes it toward risk ($+H_\tau$). For Norwegian parameters with moderate subsistence ($\zeta = 100,000$ NOK per year, $r_f = 3\%$, $\tau_1 = 1\%$):

$$H = \frac{100,000}{0.02} = 5,000,000 \text{ NOK}, \quad H_\tau = 950,000 \text{ NOK}.$$

Since $H \gg H_\tau$, the HARA effect dominates for investors with significant subsistence consumption. However, for wealthy investors with low subsistence-to-wealth ratios, the two effects are of comparable magnitude and approximately cancel, restoring near-neutrality.

Figure 4 illustrates the opposing distortions. The progressive tax shield pushes the risky share *above* the CRRA benchmark (blue curve), while the HARA subsistence floor pushes it *below* (red curve). The combined effect (purple curve) lies between: the shield partially offsets the HARA distortion, but the net effect remains conservative because $H \gg H_\tau$ at Norwegian parameters. The shaded region between the HARA-only and combined curves shows the magnitude of the progressive offset—the extent to which the tax exemption “recovers” risk-taking capacity that the subsistence floor removes.

6.5 Threshold bunching and behavioural responses

The threshold at \bar{W}_1 creates a kink in the budget set: the marginal tax rate jumps from 0 to τ_1 . Standard bunching theory (Kleven, 2016) predicts that some investors will reduce their taxable wealth to just below the threshold. The empirical evidence is substantial.

Jakobsen et al. (2020) study the Danish progressive wealth tax (rates ~ 0.7 – 1.2% , thresholds near the 98th percentile) using administrative data. They find sizable long-run elasticities of taxable wealth at the top, driven by both real savings adjustments and avoidance. Their lifecycle model with “residual wealth utility” calibrates an elasticity of taxable wealth with respect to the net-of-tax rate that implies meaningful revenue consequences from threshold design.

Garbinti et al. (2024) provide a striking finding from the French ISF: *no bunching at pure tax-rate kinks*, but *large, sharp bunching at information-requirement thresholds*. This suggests that the portfolio distortion from bracket boundaries may be smaller than the evasion response triggered by exemption thresholds—a distinction with direct implications for Norwegian threshold design.

Londoño-Vélez and Ávila Mahecha (2024) document clear bracket bunching in the Colombian wealth tax, with two-fifths of the wealthiest 0.01% hiding a third of their wealth offshore. The

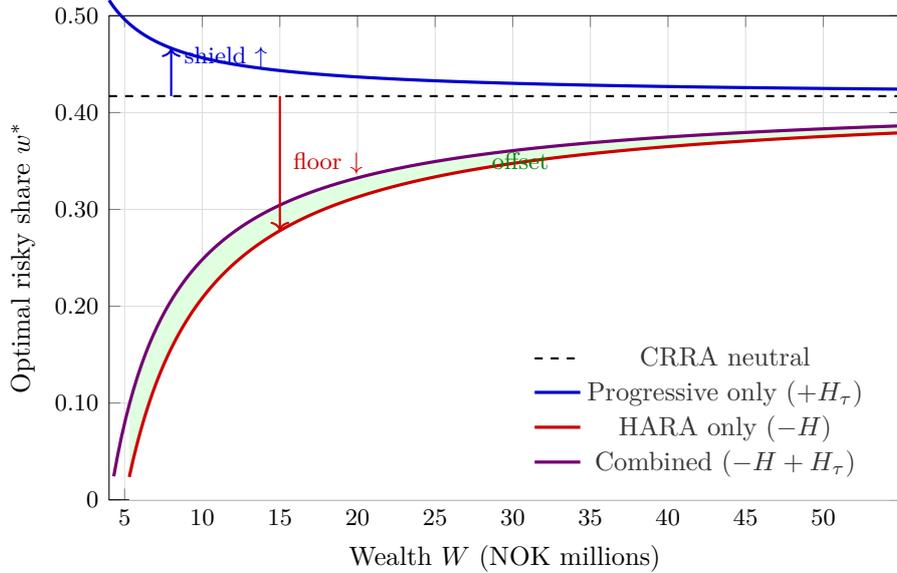


Figure 4: Opposing portfolio distortions under progressive taxation with HARA preferences. The dashed line is the CRRA benchmark. The blue curve shows the progressive-only effect (CRRA investor with threshold): the tax shield H_τ increases risk-taking. The red curve shows the HARA-only effect (proportional tax): subsistence floor H reduces risk-taking. The purple curve combines both: $w^* = w_{\text{neutral}}^* \cdot (W - H + H_\tau)/W$. The shaded region shows the magnitude of the progressive offset. Norwegian bracket 1 calibration: $H = 5\text{M NOK}$, $H_\tau = 950,000\text{ NOK}$, same parameters as Figure 1.

much larger behavioural response in Colombia compared to Scandinavia underscores the role of third-party reporting and enforcement in determining the practical relevance of threshold effects.

In Norway, the combination of comprehensive third-party reporting (for listed securities, bank deposits, and registered real estate) with self-reporting for unlisted businesses creates a dual regime. The valuation discounts analysed in Section 4 interact with the threshold: an investor can reduce taxable wealth below \bar{W}_1 by holding assets with low assessment fractions α_i , combining the non-uniform and progressive channels of non-neutrality.

6.6 Policy implications

The analysis yields three insights for threshold design.

First, a higher threshold reduces portfolio distortions for moderate-wealth investors—precisely those most affected by the progressive structure. Raising \bar{W}_1 from 1.9M to, say, 5M NOK would eliminate the tax shield for investors below 5M and reduce H_τ/W for all remaining taxpayers. This comes at a revenue cost that depends on the wealth distribution; the concentrated ownership structure of Norwegian wealth suggests the revenue loss may be moderate.

Second, the progressive structure eases the liquidity constraint. Investors just above the threshold face a low average rate $\bar{\tau}(W) = \tau_1(1 - \bar{W}_1/W)$, which mitigates the consumption-saving and liquidity effects that would arise under a fully proportional tax at the same marginal rate. This is particularly relevant for investors whose wealth is concentrated in illiquid assets (real estate,

private firms) where the wealth tax may force asset sales or borrowing.

Third, the interaction between progressivity and the non-uniform assessment of Section 4 creates compound distortions. An investor near the threshold faces strong incentives to hold assets with low assessment fractions: primary housing ($\alpha = 0.25$) can reduce taxable wealth below \bar{W}_1 for investors with gross wealth well above the threshold. To the extent that this encourages owner-occupied housing over productive investment, the progressive structure amplifies the misallocation from non-uniform assessment.

7 Endogenous Labour Supply and Entrepreneurial Effort

The preceding sections treat the investor's income as exogenous. In practice, the wealth tax may alter the incentive to supply labour or exert entrepreneurial effort, creating a feedback loop between taxation, wealth accumulation, and the labour-leisure margin. This section analyses the labour supply channel and its interaction with the progressive structure of Section 6.

7.1 Wealth effects under proportional taxation

Consider an investor who chooses portfolio weight w , labour supply ℓ , and consumption c . Labour earns income $y(\ell) = \omega \ell$ at a disutility cost $v(\ell)$, where ω is the wage rate and v is increasing and convex. The wealth dynamics become

$$dW = [W(w(\mu - r_f) + r_f - \tau_w) + \omega \ell - c] dt + w W \sigma dZ. \quad (53)$$

Under CRRA preferences over consumption and separable disutility of labour, $U = u(c) - v(\ell)$ with $u(c) = c^{1-\gamma}/(1-\gamma)$, the first-order conditions give

$$v'(\ell) = \omega u'(c). \quad (54)$$

Since the wealth tax reduces lifetime wealth and hence consumption, it raises marginal utility $u'(c)$, increasing optimal labour supply.

Proposition 7 (Separability under proportional taxation). *Under CRRA preferences with separable labour disutility and a proportional wealth tax at rate τ_w :*

- (i) *The optimal portfolio weight is $w^* = (\mu - r_f)/(\gamma\sigma^2)$, independent of labour supply ℓ .*
- (ii) *The optimal labour supply satisfies $v'(\ell) = \omega u'(c^*)$, where c^* is optimal consumption. The wealth tax affects ℓ only through the income effect: it reduces c^* , raises $u'(c^*)$, and increases ℓ .*
- (iii) *The portfolio and labour decisions are decoupled.*

Proof. The proportional tax reduces all asset returns uniformly, leaving the excess return $\mu - r_f$ and the variance σ^2 unchanged. The portfolio first-order condition depends only on these, not on ℓ . The labour first-order condition (54) depends on the marginal utility of consumption, which is affected by the tax through wealth but not through the portfolio weight. \square

Remark (Analogy to a lump-sum tax). For portfolio purposes, the proportional wealth tax is equivalent to a reduction in the risk-free rate. For labour supply purposes, it acts like a lump-sum tax proportional to wealth: it reduces disposable resources without distorting the marginal return to effort.

7.2 The threshold notch and labour supply

The progressive structure of Section 6 breaks the separability. At the threshold \bar{W}_1 , the marginal tax rate jumps from 0 to τ_1 , creating a *notch* in the budget set. An investor whose wealth would place them just above \bar{W}_1 faces a discrete choice: accumulate marginally more wealth and pay τ_1 on the excess, or reduce accumulation (by working less, consuming more, or shifting into exempt assets) to remain below the threshold.

The notch induces a *substitution effect* absent under proportional taxation. Near the threshold, the effective marginal tax rate on wealth accumulation is locally infinite (a small increase in W from $\bar{W}_1 - \epsilon$ to $\bar{W}_1 + \epsilon$ triggers a discrete tax liability on the entire excess). This creates a dominated region: investors with wealth slightly above \bar{W}_1 would be better off at exactly \bar{W}_1 .

The width of the dominated region depends on the tax rate and discount rate. In the Norwegian system, the annual tax liability at $W = \bar{W}_1 + \Delta W$ is $\tau_1 \cdot \Delta W$. For $\Delta W = 100,000$ NOK, this is 1,000 NOK per year—a modest amount that limits bunching incentives for most investors. However, the present value of the perpetual tax stream $\tau_1 \cdot \Delta W / (r_f - \tau_1)$ can be substantial: at $r_f = 3\%$, it equals $\Delta W / 2$, meaning that every 100,000 NOK above the threshold costs 50,000 NOK in present-value tax.

For entrepreneurs whose business wealth places them near the threshold, the labour supply response interacts with the portfolio response: reducing effort reduces both labour income and the value of the business, potentially bringing total wealth below \bar{W}_1 . This creates a compound distortion that does not arise under either proportional taxation or exogenous income.

7.3 Entrepreneurial effort and human capital

A distinctive feature of wealth taxation—relative to income or capital gains taxation—is its treatment of returns to entrepreneurial effort. [Güvenen et al. \(2023\)](#) formalise this through heterogeneous returns on capital: high-ability entrepreneurs earn higher returns, and a wealth tax—unlike a capital income tax—does not penalise the higher return directly. The tax liability depends on the stock of wealth, not on the flow of returns, creating a “use it or lose it” incentive: the tax erodes idle or low-return capital while leaving high-return capital relatively less burdened in after-tax terms.

In their calibrated model, replacing capital income taxation with revenue-neutral wealth taxation raises average welfare by approximately 7% of consumption-equivalent, driven by reallocation of capital from low-productivity to high-productivity entrepreneurs. A subsequent paper by the same authors extends the framework to endogenous innovation effort: when entrepreneurial productivity is a choice variable, the wealth tax further incentivises effort because entrepreneurs keep more of the upside compared to capital income taxation.

This mechanism interacts with the Norwegian tax system’s treatment of intangible assets. [Bjørneby et al. \(2023\)](#) document a positive causal relationship between wealth tax liability and employment in closely held firms, using Norwegian register data and tax reforms over 2007–2017. A 100,000 NOK increase in wealth tax liability is associated with approximately 50,000 NOK in additional wage costs. The mechanism operates through two channels: an income effect (owners work harder to compensate for the tax) and a portfolio reallocation effect (intangible assets in non-traded firms are tax-exempt, incentivising owners to invest in human capital within their businesses rather than in taxed financial assets).

This finding is consistent with the non-uniform taxation framework of Section 4: intangible business assets carry an implicit assessment fraction of $\alpha = 0$, making them the most tax-favoured asset class. The wealth tax thus creates an incentive to substitute financial capital for human capital—a reallocation that may be efficiency-enhancing if entrepreneurs have comparative advantage in their own firms.

Remark (Feedback to inelastic markets). The labour supply response also feeds back into the general equilibrium price channel of Section 5. If the threshold notch (Section 7.2) reduces entrepreneurial effort and wealth accumulation, it moderates the aggregate flow ΔF and hence the price impact. Conversely, the income effect that increases effort raises wealth accumulation and amplifies the flow. The net direction depends on which effect dominates—an empirical question that the Norwegian evidence (modest positive income elasticity) suggests resolves in favour of a small amplification.

7.4 Empirical evidence

The empirical literature on wealth effects and labour supply provides context for calibrating the magnitude of these channels.

[Scheuer and Slemrod \(2021\)](#) survey cross-country evidence and report that in Norway, households increase taxable labour income by approximately 0.01 NOK per additional NOK of wealth tax—a positive but small elasticity consistent with a modest income effect. In Switzerland and Sweden, no significant earnings response to wealth taxation was found.

The lottery and inheritance literatures provide estimates of pure wealth effects. [Picchio et al. \(2018\)](#) study Dutch lottery winners and find small but statistically significant reductions in labour earnings, with winners reducing hours but rarely exiting the labour force. Inheritance studies find that women reduce labour supply by approximately 1.5 hours per week upon receiving a bequest, while men’s supply is largely unaffected. These wealth effects run in the opposite direction from the wealth tax (which *reduces* wealth and hence *increases* labour supply), but their modest magnitude suggests that the labour supply channel is unlikely to dominate portfolio and price effects.

[Straub and Werning \(2020\)](#) overturn the classical Chamley–Judd result that long-run capital taxation should be zero. They show that when the intertemporal elasticity of substitution is below one—the empirically relevant range—the optimal long-run capital tax rate is positive

and significant. This provides theoretical support for wealth taxation as part of an optimal tax system, even when labour supply is endogenous.

7.5 Policy implications

The labour supply analysis reinforces and qualifies the findings of earlier sections.

First, a proportional wealth tax distorts labour supply only through the income effect, which operates in the “benign” direction of increasing effort. This is a welfare-relevant consideration that partially offsets the deadweight loss from any portfolio distortion.

Second, the progressive structure introduces an additional substitution effect at the threshold that can reduce effort for investors near \bar{W}_1 . A higher threshold, as discussed in Section 6.6, would eliminate this notch for moderate-wealth investors and confine the substitution effect to wealthier taxpayers for whom the threshold is less binding.

Third, the Norwegian evidence that wealth taxation increases employment in closely held firms suggests that the human capital reallocation channel may quantitatively dominate the standard labour supply distortion—at least in a system with substantial asset-class exemptions. This is a second-best argument: the non-uniform assessment that distorts portfolio choice (Section 4) may simultaneously improve labour allocation by making human capital investment tax-favoured.

Fourth, the “use it or lose it” mechanism of [Guvenen et al. \(2023\)](#) implies that the efficiency cost of wealth taxation is lower than that of capital income taxation when entrepreneurial returns are heterogeneous. This is relevant for the design of the overall tax mix: the marginal efficiency cost of raising revenue through the wealth tax may be lower than commonly assumed, particularly at moderate rates.

8 Application: A Global Minimum Wealth Tax

Two recent proposals illustrate contrasting approaches to wealth taxation at the top of the distribution. [Zucman \(2024\)](#), developed for the G20 Brazilian presidency, advocates a global minimum wealth tax of 2% on net wealth exceeding approximately \$1 billion, targeting the roughly 3,000 individuals who currently pay effective tax rates of 0.2–0.3% of their wealth; the estimated revenue is \$200–250 billion annually. In France, the National Assembly adopted a domestic variant in February 2025: a minimum effective tax of 2% on net wealth exceeding €100 million, affecting approximately 1,800 households ([Assemblée Nationale, 2025](#)). Although both share the 2% rate and Zucman’s intellectual framework, they differ in threshold, scope, and mechanism in ways that activate different non-neutrality channels. This section evaluates both proposals through the lens of the extensions developed in Sections 3–7.

8.1 The global proposal: assessment through the extension framework

Non-homothetic preferences (Section 3). For billionaires, wealth vastly exceeds any plausible subsistence floor: $W \gg H$, so $H/W \approx 0$. The HARA distortion of Proposition 3 is

negligible; CRRA is an excellent approximation. The portfolio distortion from non-homothetic preferences is a non-issue for the target population.

Non-uniform assessment (Section 4). A defining design feature of the Saez–Zucman proposal is comprehensive market-value assessment: all assets are taxed at market value with no valuation discounts. This eliminates the non-uniform channel entirely. The portfolio distortion Δw^* from Proposition 5 is zero when $\alpha_i = 1$ for all i . By contrast, the Norwegian system’s extensive valuation discounts (Section 4.3) create substantial tilts toward housing and away from deposits. The Saez–Zucman design thus avoids a major source of non-neutrality present in existing wealth taxes.

Progressive structure (Section 6). The extremely high threshold (\sim \$1 billion) means the tax shield from Proposition 6 is negligible relative to wealth. At a 2% rate with $r_f = 3\%$:

$$H_\tau = \frac{0.02 \times 10^9}{0.03 - 0.02} = 2 \times 10^9,$$

so $H_\tau/W = 2$ for $W = \$1\text{B}$ (the threshold investor) but $H_\tau/W = 0.02$ for $W = \$100\text{B}$. While the threshold effect is mechanically large at the boundary, the target population is concentrated well above it. Moreover, the proposal is designed as a *minimum* tax—a top-up on existing obligations—so the effective threshold is even higher for investors already paying wealth or capital taxes.

Inelastic markets (Section 5). This is the channel most relevant to the proposal. Billionaires hold concentrated equity positions, often representing controlling stakes in publicly traded firms. As a conservative upper bound (assuming the full liability is met by equity sales; see Section 5.4 for the payment-hierarchy refinement), the aggregate tax-induced flow is

$$\frac{\Delta F}{P_{\text{eq}}} = \tau_w \cdot \phi_B,$$

where ϕ_B is the fraction of global equity market capitalisation held by the $\sim 3,000$ affected individuals. Global billionaire wealth is approximately \$14 trillion (2024), of which roughly 60% is in equities. Against a global equity market capitalisation of \sim \$110 trillion, $\phi_B \approx 8.4/110 \approx 0.08$. The annual flow-to-capitalisation ratio is

$$\frac{\Delta F}{P_{\text{eq}}} = 0.02 \times 0.08 = 0.0016 = 0.16\%.$$

Under the inelastic markets hypothesis with $M = 5$, the annual price impact is

$$\frac{\Delta P}{P} = -M \cdot 0.16\% = -0.8\%.$$

This is a non-trivial but modest annual effect. Compounded over a holding period and borne by *all* equity holders (not only the taxed billionaires), it represents a general equilibrium cost of the proposal that the partial-equilibrium revenue estimate does not capture.

Two caveats apply. First, the estimate assumes that billionaires liquidate equities to pay the tax, whereas in practice they may use dividends, borrow against holdings, or sell other assets. Second, the multiplier M is estimated for the US equity market; the effective multiplier for the concentrated, often illiquid holdings typical of billionaire portfolios may differ substantially.

Labour supply (Section 7). The income effect of a 2% annual tax on multi-billion-dollar wealth is large in absolute terms (\$20M per year on \$1B) but small relative to the wealth base. The “use it or lose it” mechanism is directly relevant: a 2% annual tax erodes the wealth of passive heirs ($\sim 20\%$ per decade in real terms) while imposing a lower effective burden on high-return entrepreneurs. This selection effect is a feature, not a bug, of the proposal: it shifts wealth from low-return to high-return holders over time.

8.2 The French variant: a national minimum wealth tax

The French proposal shares the 2% rate and minimum-tax logic of the global blueprint but differs in three dimensions that alter the non-neutrality profile.

Lower threshold (€100M vs \$1B). The progressive channel is substantially more relevant. At the French threshold with $r_f = 3\%$:

$$H_\tau = \frac{0.02 \times 10^8}{0.03 - 0.02} = 2 \times 10^8 = \text{€}200\text{M},$$

so $H_\tau/W = 2$ for the threshold investor ($W = \text{€}100\text{M}$) and $H_\tau/W = 0.2$ even at $W = \text{€}1\text{B}$. By Proposition 6, the tax shield increases risk-taking for investors near the threshold, and the effect remains quantitatively important across a larger fraction of the affected population than under the global proposal. Combined with the HARA channel—still small at this wealth level ($H/W < 0.01$) but non-negligible in percentage terms—the net distortion is larger than for billionaires.

Minimum tax mechanism. The French proposal is explicitly a top-up: investors already paying at least 2% of their wealth through income, capital gains, and property taxes owe nothing additional. This creates a complementarity between wealth taxation and income realisation that is absent from a pure wealth tax. In our framework, the effective wealth tax rate is

$$\tau_w^{\text{eff}}(W, Y) = \max(0, 0.02 - T_{\text{existing}}(Y)/W),$$

where $T_{\text{existing}}(Y)$ is total existing tax paid and Y denotes realised income. This has two implications for portfolio choice. First, investors with high-dividend or high-turnover portfolios face a lower effective wealth tax, which reduces the non-uniform assessment channel for yield-generating assets. Second, the top-up structure creates an incentive to realise income up to the 2% threshold, reversing the usual lock-in effect of capital gains taxation—a behavioural response outside the scope of our static framework but worth noting.

Scope, portfolio geography, and market impact. The tax applies to the *global* assets of French tax residents, not only to French-domiciled holdings. With $\sim 1,800$ affected households, the aggregate wealth subject to the tax is approximately €400–500 billion, of which roughly 60% is in equities (\sim €270 billion). These equity holdings are internationally diversified: French ultra-high-net-worth portfolios typically have substantial exposure to US, European, and emerging markets, with a home-bias share in French-listed equities that we denote h . The tax-induced selling flow is therefore split across multiple markets.

Against global equity capitalisation of \sim €100 trillion, the aggregate flow ratio is modest:

$$\frac{\Delta F_{\text{global}}}{P_{\text{eq}}^{\text{global}}} = \frac{0.02 \times 270}{100,000} \approx 0.005\%,$$

an order of magnitude smaller than the global Saez–Zucman proposal. However, with home bias the flow is disproportionately concentrated in French equities. If $h = 0.4$ (a plausible estimate for concentrated family holdings), the Euronext-specific flow ratio is

$$\frac{\Delta F_{\text{FR}}}{P_{\text{eq}}^{\text{FR}}} = \frac{0.02 \times 0.4 \times 270}{3,500} \approx 0.06\%,$$

where €3.5 trillion is the Euronext Paris capitalisation. With a multiplier $M = 5$, the local price impact is $\Delta P/P \approx -0.3\%$ —smaller than the -0.8% global estimate, reflecting the dilution of selling pressure across international markets. For individual stocks in which affected taxpayers hold controlling stakes, the local impact could be substantially larger, particularly if stakes are illiquid and represent a high fraction of the free float.

The Brühlhart et al. (2022) migration channel is also more salient for a national tax: the 24% inter-cantonal response they estimate for the Swiss wealth tax is likely a lower bound for cross-border migration within the EU’s single market, absent coordinated enforcement. The combination of global asset coverage with national jurisdiction creates a tension: the tax base is mobile even if the assets themselves are not.

8.3 Comparison with existing wealth taxes

Table 4 summarises how the extension channels apply to the Norwegian system, the global Saez–Zucman proposal, and the French variant.

The comparison reveals a spectrum of design trade-offs. The Norwegian system imposes a low rate (1.0–1.1%) on a broad population but introduces substantial non-neutralities through valuation discounts, progressive thresholds, and debt deductibility. The global Saez–Zucman proposal imposes a higher rate (2%) on a very narrow population and eliminates most non-neutrality channels through comprehensive market-value assessment and an extremely high threshold; the residual distortion is concentrated in the inelastic markets channel. The French variant occupies an intermediate position: its lower threshold activates the progressive channel for a larger fraction of affected taxpayers, while the top-up mechanism partially mitigates the non-uniform assessment channel by crediting existing taxes against the wealth tax liability. All three designs share an irreducible non-neutrality from the inelastic markets channel—an unavoidable

Table 4: Non-neutrality channels under three wealth tax designs.

Channel	Norwegian	Global (Zucman)	French variant
Non-homothetic (HARA)	Moderate ($H/W \sim 0.1\text{--}0.5$)	Negligible ($H/W \approx 0$)	Small ($H/W < 0.01$)
Non-uniform assessment	Large (tilts to housing)	Zero (market values)	Reduced (top-up offsets)
Progressive threshold	Significant (near NOK 1.9M)	Negligible (threshold \sim \\$1B)	Significant (threshold €100M)
Inelastic markets	Small, local ($\phi \approx 0.2$)	Moderate, global ($\phi_B \approx 0.08$)	Small–moderate (diluted globally)
Labour supply	Positive (human capital)	“Use it or lose it” (selection effect)	“Use it or lose it” + income realisation

consequence of any tax that induces portfolio rebalancing in a market with finite liquidity.

8.4 Implementation challenges

The practical feasibility of the proposal depends on two conditions outside our model.

First, *international coordination*. Without universal adoption, the proposal creates incentives for tax-base migration. Brülhart et al. (2022) estimate that 24% of the Swiss wealth tax response comes from inter-cantonal migration; at the global level, the relevant margin is jurisdiction shopping among sovereign states. Londoño-Vélez and Ávila Mahecha (2024) document that two-fifths of the wealthiest Colombians hide a third of their wealth offshore, underscoring the enforcement challenge.

Second, *valuation of illiquid assets*. The proposal requires annual market-value assessment of private companies, art, real estate, and other assets that do not trade on liquid markets. This is precisely the valuation problem that drives the non-uniform assessment of Section 4: Norwegian valuation discounts exist in part because market values are difficult to determine. For listed equities and bank deposits, the problem is solved by third-party reporting; for private wealth, it remains the binding constraint on implementation.

9 Tax-Induced Migration

The preceding sections analyse how investors respond to wealth taxation by adjusting portfolios, labour supply, and asset prices. A more radical response is to change tax jurisdiction entirely. Under European-style domicile-based taxation, a change of residence is sufficient to eliminate the wealth tax liability; the progressive threshold of Section 6 then implies a *participation margin* at which investors exit the tax base altogether.

9.1 Migration as a threshold response

Consider an investor with wealth $W > \bar{W}_1$ (the first taxable threshold) who faces a heterogeneous migration cost $c_i > 0$, reflecting the disutility of relocation: social ties, language, business

location-specificity, family considerations, and regulatory frictions. Under domicile-based taxation, the investor migrates if the present value of future wealth tax liabilities exceeds the migration cost.

Within the progressive framework of Section 6, the present value of the tax flow for an investor in bracket j is

$$\text{PV}_\tau(W) = \frac{\tau_j W - R_j}{r_f - g + \tau_j}, \quad (55)$$

where g is the expected real growth rate of wealth and the denominator reflects that the tax base grows at rate $g - \tau_j$ (net of wealth erosion from the tax itself). The investor migrates when $\text{PV}_\tau(W) > c_i$.

This yields a migration threshold

$$W_i^* = \frac{c_i(r_f - g + \tau_j) + R_j}{\tau_j}. \quad (56)$$

Investors with $W > W_i^*$ migrate; those with $W < W_i^*$ stay. Several features are immediate.

First, the migration threshold is *decreasing* in τ_j : higher wealth tax rates lower the threshold at which migration becomes attractive, expanding the set of investors who leave. This is the intensive margin that drives the empirical relationship between wealth tax rates and out-migration.

Second, the progressive rebate R_j *raises* the migration threshold, since part of the tax liability is offset by the cumulative exemption. A higher basic exemption (\bar{W}_1) therefore has a dual effect: it eliminates the participation margin for investors near the threshold (Section 6) *and* raises the migration cost–benefit ratio for investors above it.

Third, under a citizenship-based tax system—as in the United States—relocation does not eliminate the wealth tax liability. The effective migration cost becomes $c_i + \text{PV}_{\text{exit}}$, where PV_{exit} includes the present value of continued tax obligations and any exit tax on unrealised gains. This structural difference explains the finding of Young et al. (2016) that US millionaire migration is low (2.4% annually, below the general population rate): for American taxpayers, physical relocation alone is insufficient to escape the tax base. This result is therefore not transferable to the European context, where domicile-based taxation makes migration a complete opt-out.

Figure 5 illustrates the mechanism in (W, c_i) space. Each line represents the indifference locus $c_i = \text{PV}_\tau(W)$ for a given tax rate: investors below the line (low migration cost relative to tax burden) migrate; those above stay. All lines emanate from the exemption threshold $(\bar{W}_1, 0)$, since at the threshold the tax liability is zero and no migration incentive exists regardless of rate. A tax increase from 0.85% to 1.0% rotates the line upward, expanding the “migrate” region: the shaded wedge between the two lines represents investors who were previously in the “stay” region but are pushed into migration by the rate increase. At higher rates (2%, the Zucman proposal), the “migrate” region expands substantially, underscoring the coordination argument of Section 8.

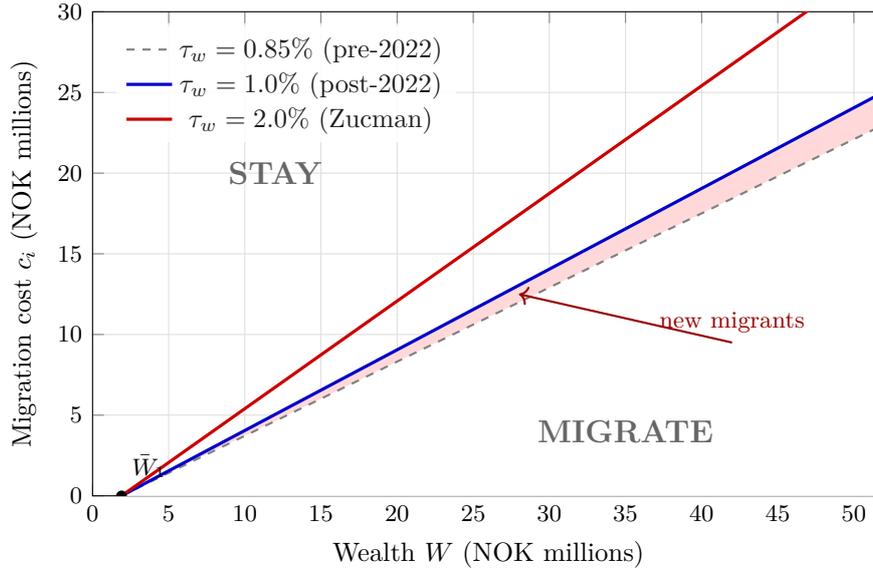


Figure 5: Migration decision in (W, c_i) space. Each line shows the indifference locus $c_i = PV_\tau(W)$: investors below the line migrate; those above stay. A tax increase from 0.85% to 1.0% rotates the line upward, expanding the “migrate” region (shaded wedge). The 2% Zucman rate further expands it. All lines emanate from the exemption threshold $\bar{W}_1 = 1.9\text{M NOK}$. Under citizenship-based taxation (US), the effective cost shifts upward by PV_{exit} , compressing the “migrate” region. Calibration: $r_f = 3\%$, $g = 2\%$.

9.2 The Norwegian case: 2022–2024

Norway provides a contemporary case study that illustrates both the migration response and the identification challenges. The Støre government increased the effective wealth tax burden through several simultaneous channels beginning in 2022: the wealth tax rate rose from 0.85% to 1.1%, the effective tax rate on dividends and capital gains increased from approximately 31.7% to 37.8% (through the *oppjusteringsfaktor*, the multiplicative upward adjustment applied before the 22% ordinary income tax rate), and the five-year expiration period for the exit tax on unrealised gains was abolished in November 2022.

The behavioural response was swift. According to the Norwegian Ministry of Finance, 82 high-wealth individuals with combined net wealth of approximately NOK 46 billion relocated in 2022–2023, more than the previous 13 years combined. The Civita think tank reports that 261 residents with wealth above NOK 10 million left in 2022, roughly double the pre-reform baseline. Switzerland—which has low cantonal wealth taxes, no capital gains tax, and established Norwegian expatriate communities—was the primary destination.

However, the identification of a *wealth tax* effect is confounded by at least three concurrent factors.

Confounding with dividend and capital gains taxation. The effective dividend tax rate increased by approximately 6 percentage points simultaneously with the wealth tax increase. For business owners who extract dividends to cover wealth tax liabilities—a common pattern documented by Bjørneby et al. (2023)—the two taxes interact: the wealth tax creates a liquidity

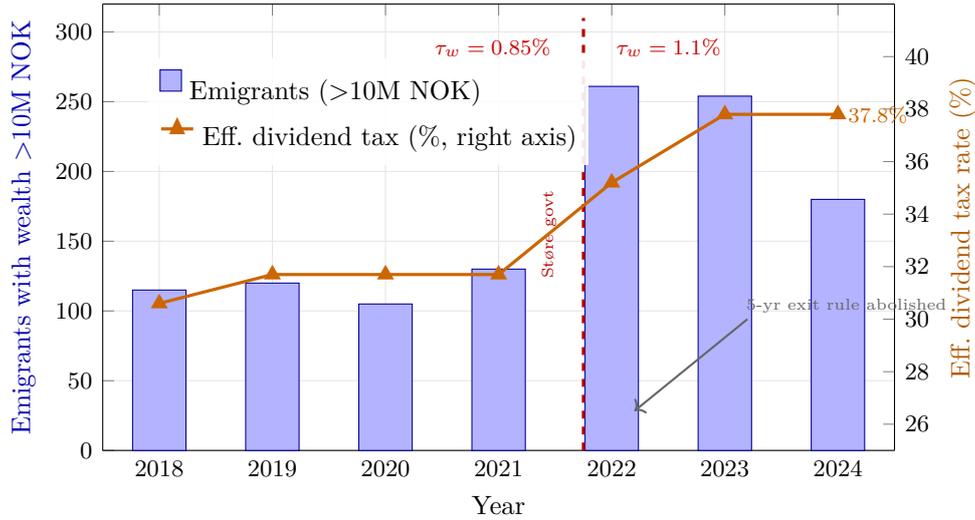


Figure 6: Norwegian high-wealth emigration and concurrent tax changes, 2018–2024. Bars show annual emigrants with wealth above NOK 10 million (left axis; Civita estimates; 2024 figure is preliminary). The line shows the effective dividend tax rate (right axis). The wealth tax rate is annotated at the reform boundary ($0.85\% \rightarrow 1.1\%$). The emigration spike in 2022 coincides with increases in both rates and the abolition of the five-year exit tax expiration (arrow), illustrating the identification challenge.

need that triggers dividend extraction, which is itself taxed at the higher rate. Separating the marginal contribution of each tax requires variation in one holding the other constant, which the 2022 reform does not provide.

Confounding with the exit tax. The abolition of the five-year exit tax expiration in November 2022 changed the calculus for investors with large unrealised capital gains. Under the old regime, emigration followed by a five-year holding period eliminated both the wealth tax and the capital gains tax on accumulated gains—a powerful joint incentive. Some high-profile departures occurred just before the November deadline, suggesting that the exit tax change, rather than the wealth tax increase alone, was the binding margin for investors with concentrated equity positions.

Confounding with political and structural factors. The change from a conservative to a centre-left government in 2021 brought broader shifts in economic policy, regulatory approach, and political rhetoric. Wealthy individuals may respond not only to enacted tax changes but also to perceived policy direction: the expectation of further increases, proposed regulations on private equity, and public debate about wealth concentration may jointly affect the relocation decision. Disentangling these channels from the tax rate itself requires careful empirical work beyond the scope of a theoretical framework.

Figure 6 summarises the identification challenge visually: the emigration spike in 2022 coincides with simultaneous changes in the wealth tax rate, the effective dividend tax rate, and the exit tax regime. No single-variable causal attribution is possible from the time series alone.

9.3 Decomposing behavioural responses

The migration margin interacts with other behavioural channels. Brülhart et al. (2022) decompose the Swiss wealth tax response into three components: taxpayer migration (24%), capitalisation into housing prices (21%), and evasion or avoidance (55%). Seim (2017) finds that approximately one-third of the Swedish response reflects underreporting. In the Norwegian context, Iacono and Smedsvik (2024) exploit the Bø municipality experiment—in which a single municipality unilaterally reduced its municipal wealth tax rate from 0.70% to 0.20% in 2021—and find that taxpayer mobility accounts for 79% of the increase in taxable wealth.

These decompositions suggest a hierarchy of responses. Evasion and avoidance (including legal restructuring) are the first margin, as they carry lower personal costs than physical relocation. Migration is the margin of last resort—activated when the tax burden is sufficiently large relative to migration costs, consistent with the threshold model of equation (56). The Agrawal et al. (2025) finding that Spanish mobility responds to large discrete tax differentials (Madrid as a regional tax haven) but not to small inter-regional variation supports this interpretation: migration requires a sufficiently large $PV_\tau(W) - c_i$ gap.

9.4 Implications for the framework

The migration channel has two implications for the non-neutrality analysis.

First, it introduces a *selection effect* on the remaining tax base. If the highest-wealth investors migrate, the population remaining in the tax jurisdiction is truncated from above. This changes the composition of the investor base in the inelastic markets analysis (Section 5): the investors who remain may hold less concentrated, more diversified portfolios, potentially reducing the aggregate demand multiplier. Paradoxically, Norway’s wealth tax revenues *increased* from NOK 27 billion (2022) to a projected NOK 34 billion (2025) despite the exodus, because the broader middle-wealth tax base remained intact and asset values appreciated.

Second, the threat of migration creates an implicit constraint on the wealth tax rate. In a Tiebout framework, jurisdictional competition for mobile tax bases limits the feasible tax rate to the level at which the marginal migrant is indifferent. This is precisely the coordination argument for the Saez–Zucman global proposal (Section 8): a minimum tax enforced across jurisdictions eliminates the migration margin by removing the outside option. The failure of unilateral national proposals—such as the French variant of Section 8.2, rejected in part due to concerns about capital flight—illustrates the coordination constraint.

A full welfare analysis of the migration channel requires an empirical estimate of the migration cost distribution $F(c)$, which determines both the revenue loss from emigration and the dead-weight loss from distorted location decisions. Kleven et al. (2020) emphasise that migration elasticities are not structural parameters: they depend on population characteristics, jurisdiction size, the degree of international tax coordination, and the availability of evasion alternatives. This is fundamentally an empirical identification problem that our theoretical framework can frame but not resolve.

10 Conclusion

This paper has investigated the robustness of wealth tax neutrality along two dimensions: the breadth of conditions under which neutrality holds, and the channels through which it breaks.

On the first dimension, we have shown that portfolio neutrality extends well beyond the geometric Brownian motion and location-scale settings of Frøseth (2026). Under CRRA preferences, neutrality is preserved in stochastic volatility models—including the Heston model and general Markov diffusions—where it applies to all intertemporal hedging demands, not just the myopic component (Propositions 1 and 2). It is also preserved under Epstein–Zin recursive utility, where the separation of risk aversion from the elasticity of intertemporal substitution does not affect the homogeneity that drives the result (Proposition 4). The operative condition remains CRRA: non-homothetic preferences such as HARA break neutrality by making the effective risk aversion wealth-dependent (Proposition 3).

On the second dimension, we have identified four channels through which implemented wealth taxes depart from neutrality, even when investors have CRRA preferences. Non-uniform assessment across asset classes tilts portfolios toward assets with lower valuation fractions, a distortion we calibrate to the Norwegian system (Proposition 5). General equilibrium effects in inelastic markets amplify fundamental price adjustments through a demand multiplier, creating a wedge between partial and general equilibrium outcomes (Section 5). Progressive threshold structures introduce a tax shield that increases risk-taking for investors near the exemption threshold, an effect that operates in the opposite direction from HARA and partially offsets it when both are present (Proposition 6). Endogenous labour supply is separable from portfolio choice under proportional taxation but introduces additional distortions at progressive thresholds (Proposition 7). At the extreme, the progressive threshold generates a participation margin: tax-induced migration (Section 9). Under European domicile-based taxation, relocation eliminates the entire wealth tax liability once the present value of that liability exceeds the investor’s migration cost. The Norwegian experience after 2022 illustrates both the reality of this response and the difficulty of isolating it from concurrent changes to dividend taxation, exit tax rules, and the political environment. The application to the Saez–Zucman global proposal and its French national variant (Section 8) illustrates how design choices—threshold level, scope, and coordination—activate different subsets of these channels.

Directions for further work

Several extensions remain open.

Theoretical. The multi-period consumption-saving problem with endogenous labour supply has been studied in separate literatures but not in the presence of a wealth tax. A model that jointly determines consumption, portfolio choice, and labour supply under progressive wealth taxation would unify the channels analysed here. General equilibrium with heterogeneous agents—in particular, different wealth tax rates across brackets—would capture the redistributive dynamics that a representative-agent framework cannot. The interaction between wealth tax, capital gains tax, and income tax in a joint optimal design problem is a natural next

step, though the presence of foreign investors subject to different tax regimes adds substantial complexity. Finally, extending the security market fan of Frøseth (2026) to multi-factor models would clarify how the wealth tax interacts with factor premia.

Empirical. The non-uniform assessment channel can be tested directly using Norwegian wealth registry data: the introduction and removal of valuation discounts provide natural experiments for estimating portfolio reallocation. The progressive threshold channel predicts bunching behaviour near the exemption boundary, which can be tested using the administrative data employed by Jakobsen et al. (2020) for Denmark and Garbinti et al. (2024) for France. The inelastic markets channel yields testable predictions about asset price responses to wealth tax reforms: Brühlhart et al. (2022) provide a template using Swiss cantonal variation, while the Norwegian system offers variation across asset classes with different assessment fractions. On the labour supply side, the “use it or lose it” mechanism of Guvenen et al. (2023) can be tested using the relationship between entrepreneurial returns and wealth tax payments in the Norwegian firm register. The migration channel (Section 9) poses the hardest identification challenge: Jakobsen et al. (2024) estimate that a one percentage point increase in the top wealth tax rate reduces the stock of wealthy taxpayers by approximately 2%, but as Kleven et al. (2020) emphasise, migration elasticities are not structural parameters but depend on population characteristics, jurisdiction size, the degree of international tax coordination, and the availability of evasion alternatives. The Norwegian post-2022 experience, where wealth tax, dividend tax, and exit tax changes occurred simultaneously alongside a change in political direction, illustrates why isolating the wealth tax margin requires careful empirical design—ideally exploiting variation in one instrument while holding the others constant. Finally, the empirical literature on wealth taxation and firm behaviour is dominated by traditional private firms; evidence on intangible-heavy, venture-capital-backed firms—where book-value assessment creates the largest effective discount and liquidity frictions are most acute—remains scarce, despite these firms’ prominence in the current policy debate.

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.

A Leverage and debt deductibility: historical examples

This appendix provides calibrated examples of the sheltering strategy discussed in Section 4.4.

Under the system prevailing before Norway’s proportional debt reduction rules, debt was deductible at face value ($\beta_i = 1$) while assets received valuation discounts. Consider commercial property under the old 45% discount ($\alpha = 0.55$) financed at 80% loan-to-value ($\ell = 0.80$):

$$\frac{\text{Taxable wealth}}{V} = 0.55 - 0.80 = -0.25.$$

Each krone of leveraged commercial property *reduced* the investor’s tax base by 0.25 kroner. An investor holding NOK 10 million in bank deposits (assessed at full value) could eliminate the associated tax liability entirely by accumulating approximately NOK 40 million of leveraged commercial property, whose negative taxable contribution offsets the deposits. Once the tax base reaches zero (its floor), further leveraged property yields no additional tax benefit, creating a kink in the effective marginal incentive.

For primary housing ($\alpha = 0.25$) with moderate leverage ($\ell = 0.60$), the sheltering capacity was even larger: $0.25 - 0.60 = -0.35$ per krone of asset value.

Norway’s proportional debt reduction rules now set $\beta_i = \alpha_i$ for discounted assets, eliminating negative tax bases. For shares and commercial property with a 20% discount, the associated debt deduction is reduced by 20%, giving $\beta_i = 0.80$. While this eliminates the most egregious sheltering, the effective tax rate on equity in discounted asset classes ($\tau_w \alpha_i$) remains below the statutory rate, preserving a residual incentive to hold leveraged positions in discounted assets.

B Time-scale dependence of stylised facts

This appendix provides the detailed discussion of how the stylised facts motivating the stochastic volatility extension vary across time horizons, supplementing the summary in Section 2.8.

Among Cont’s (2001) eleven stylised facts, the most dramatic departures from GBM are high-frequency phenomena. The inverse cubic power law for return tails (Gopikrishnan et al., 1999; Gabaix et al., 2003) attenuates at longer horizons via aggregational Gaussianity, with a crossover to approximately Gaussian tails at horizons of several weeks to a few months (Bouchaud and Potters, 2003). Volatility clustering weakens from long-range dependence at daily frequency (Ding et al., 1993) to short-memory persistence at monthly and annual horizons, where the Heston model’s exponential autocorrelation is a reasonable approximation.

What *does* persist at policy-relevant horizons is time-varying expected returns (Campbell and Shiller, 1988), regime-like volatility dynamics, and the variance risk premium—precisely the features captured by the general Markov diffusion framework of Section 2.6. The stochastic volatility extension is therefore valuable not because high-frequency anomalies require it, but because it demonstrates that neutrality extends to the empirically relevant features that *do* persist: time-varying investment opportunities, hedging demands, and stochastic risk premia. This also diminishes the urgency of extending the result to jump-diffusion processes, since the power-law tails that jumps would capture are primarily a high-frequency phenomenon.

C Empirical evidence on non-uniform assessment

This appendix reviews the empirical literature on wealth tax effects under non-uniform assessment, supplementing the summary in Section 4.5.

Ring (2024) provides the most direct empirical test of wealth tax neutrality under non-uniform assessment. Exploiting geographic discontinuities in housing valuations created by Norway’s

2010 reform, Ring estimates the causal effect of the wealth tax on household saving and portfolio allocation. His central finding is that the wealth tax has a positive effect on financial saving (elasticity of approximately 3.76), as households save to offset the tax, but the portfolio *composition* between risky and safe assets is unaffected when the tax does not discriminate between them. This is consistent with the neutrality result under uniform assessment established in the main text.

However, when the assessment differentials create distinct after-tax returns, portfolio distortions emerge. Fagereng et al. (2024) study how investors respond to changes in the equity premium induced by wealth tax reforms. They find that households do adjust their risky asset shares in response to differential tax treatment, but slowly: full reoptimisation takes approximately five years. This sluggish adjustment is consistent with portfolio inertia and suggests that the theoretical distortions derived in Section 4 represent a long-run equilibrium that is approached gradually.

Brühlhart et al. (2022) study wealth tax responses in Switzerland, where cantonal rate variation creates quasi-experimental conditions. They find that a one percentage point reduction in the wealth tax rate increases reported taxable wealth by approximately 43%, though much of this response reflects taxpayer mobility between cantons and valuation responses rather than real portfolio reallocation.

References

- David R. Agrawal, Dirk Foremny, and Clara Martínez-Toledano. Wealth tax mobility and tax coordination. *American Economic Journal: Applied Economics*, 17(1):402–430, 2025. doi: 10.1257/app.20220615.
- Assemblée Nationale. Proposition de loi instaurant une contribution minimale sur le patrimoine des ultra-riches, 2025. Adopted by the National Assembly on 20 February 2025 (116–39); rejected by the Senate, 12 June 2025.
- Gurdip S. Bakshi and Zhiwu Chen. The spirit of capitalism and stock-market prices. *American Economic Review*, 86(1):133–157, 1996.
- Janis Berzins, Øyvind Bøhren, and Bogdan Stacescu. Shareholder illiquidity and firm behavior: Financial and real effects of the personal wealth tax. *Review of Finance*, 26(5):1165–1203, 2022.
- 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.
- 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 Marc Potters. *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management*. Cambridge University Press, 2nd edition, 2003. doi: 10.1017/CBO9780511753893.

- Jean-Philippe Bouchaud, J. Doyne Farmer, and Fabrizio Lillo. How markets slowly digest changes in supply and demand. In Thorsten Hens and Klaus Reiner Schenk-Hoppé, editors, *Handbook of Financial Markets: Dynamics and Evolution*, pages 57–160. North-Holland, 2009. doi: 10.1016/B978-012374258-2.50006-3.
- 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.
- Markus K. Brunnermeier and Stefan Nagel. Do wealth fluctuations generate time-varying risk aversion? Micro-evidence on individuals’ asset allocation. *American Economic Review*, 98(3): 713–736, 2008. doi: 10.1257/aer.98.3.713.
- Laurent E. Calvet and Paolo Sodini. Twin picks: Disentangling the determinants of risk-taking in household portfolios. *The Journal of Finance*, 69(2):867–906, 2014. doi: 10.1111/jofi.12125.
- John Y. Campbell and Robert J. Shiller. Stock prices, earnings, and expected dividends. *The Journal of Finance*, 43(3):661–676, 1988. doi: 10.1111/j.1540-6261.1988.tb04598.x.
- Christopher D. Carroll. Why do the rich save so much? In Joel B. Slemrod, editor, *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, pages 465–484. Harvard University Press, 2000. NBER Working Paper No. 6549, 1998.
- 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.
- Rama Cont. Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001. doi: 10.1080/713665670.
- Joshua Coval and Erik Stafford. Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics*, 86(2):479–512, 2007. doi: 10.1016/j.jfineco.2006.09.007.
- Zhuanxin Ding, Clive W. J. Granger, and Robert F. Engle. A long memory property of stock market returns and a new model. *Journal of Empirical Finance*, 1(1):83–106, 1993. doi: 10.1016/0927-5398(93)90006-D.
- Adrian A. Drăgulescu and Victor M. Yakovenko. Probability distribution of returns in the Heston model with stochastic volatility. *Quantitative Finance*, 2(6):443–453, 2002. doi: 10.1088/1469-7688/2/6/303.
- Larry G. Epstein and Stanley E. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969, 1989. doi: 10.2307/1913778.
- Andreas Fagereng, Luigi Guiso, and Marius A. K. Ring. How much and how fast do investors respond to equity premium changes? evidence from wealth taxation. Working paper, 2024.

- Anders G. Frøseth. Asset returns, portfolio choice, and proportional wealth taxation. 2026. arXiv:2603.05264 [physics.soc-ph].
- Xavier Gabaix and Ralph S. J. Koijen. In search of the origins of financial fluctuations: The inelastic markets hypothesis. NBER Working Paper No. 28967, revised 2023, 2021.
- Xavier Gabaix, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley. A theory of power-law distributions in financial market fluctuations. *Nature*, 423:267–270, 2003. doi: 10.1038/nature01624.
- Bertrand Garbinti, Jonathan Goupille-Lebret, Thomas Piketty, Clara Martínez-Toledano, and Gabriel Zucman. Tax design, information, and elasticities: Evidence from the French wealth tax. World Inequality Lab Working Paper, 2024.
- Parameswaran Gopikrishnan, Martin Meyer, Luís A. Nunes Amaral, and H. Eugene Stanley. Inverse cubic law for the distribution of stock price variations. *The European Physical Journal B*, 3(2):139–140, 1999. doi: 10.1007/s100510050292.
- Fatih Guvenen, Gueorgui Kambourov, Burhanettin Kuruscu, Sergio Ocampo-Diaz, and Daphne Chen. Use it or lose it: Efficiency and redistributive effects of wealth taxation. *The Quarterly Journal of Economics*, 138(2):835–894, 2023. doi: 10.1093/qje/qjac047.
- 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.
- Roberto Iacono and Bård Smedsvik. Behavioral responses to wealth taxation: Evidence from a Norwegian reform. WID World Working Paper 2023/30, revised April 2024, 2024.
- 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, 2024.
- Henrik Kleven, Camille Landais, Mathilde Muñoz, and Stefanie Stantcheva. Taxation and migration: Evidence and policy implications. *Journal of Economic Perspectives*, 34(2):119–142, 2020. doi: 10.1257/jep.34.2.119.
- Henrik J. Kleven. Bunching. *Annual Review of Economics*, 8:435–464, 2016. doi: 10.1146/annurev-economics-080315-015234.
- Ralph S. J. Koijen and Motohiro Yogo. A demand system approach to asset pricing. *Journal of Political Economy*, 127(4):1475–1515, 2019. doi: 10.1086/701683.
- Juliana Londoño-Vélez and Javier Ávila Mahecha. Behavioral responses to wealth taxation:

- Evidence from Colombia. *The Review of Economic Studies*, 92(4):2624–2668, 2024. doi: 10.1093/restud/rdae052.
- Robert C. Merton. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413, 1971. doi: 10.1016/0022-0531(71)90038-X.
- Robert C. Merton. An intertemporal capital asset pricing model. *Econometrica*, 41(5):867–887, 1973. doi: 10.2307/1913811.
- OECD. *The Role and Design of Net Wealth Taxes in the OECD*. Number 26 in OECD Tax Policy Studies. OECD Publishing, Paris, 2018. doi: 10.1787/9789264290303-en.
- José Luis Peydró, Hernán Rincón-Castro, Miguel Sarmiento-Paipilla, and Alejandro Granados. Wealth taxes and firms’ capital structures: Credit supply and real effects. Borradores de Economía 1316, Banco de la República de Colombia, 2025.
- Matteo Picchio, Sigrid Suetens, and Jan C. van Ours. Labour supply effects of winning a lottery. *The Economic Journal*, 128(611):1700–1729, 2018. doi: 10.1111/eoj.12465.
- Marius A. K. Ring. Wealth taxation and household saving: Evidence from assessment discontinuities in Norway. *The Review of Economic Studies*, 92(5):3375–3411, 2024. doi: 10.1093/restud/rdae091.
- Florian Scheuer and Joel Slemrod. Taxing our wealth. *Journal of Economic Perspectives*, 35(1):207–230, 2021. doi: 10.1257/jep.35.1.207.
- David Seim. Behavioral responses to wealth taxes: Evidence from Sweden. *American Economic Journal: Economic Policy*, 9(4):395–421, 2017. doi: 10.1257/pol.20150290.
- Skatteetaten. Verdssettingsrabatt ved fastsetting av formue. <https://www.skatteetaten.no/en/person/taxes/get-the-taxes-right/valuation-discount-in-connection-with-assessment-of-wealth/>, 2026. Accessed February 2026.
- Ludwig Straub and Iván Werning. Positive long-run capital taxation: Chamley–Judd revisited. *American Economic Review*, 110(1):86–119, 2020. doi: 10.1257/aer.20150210.
- Jessica A. Wachter and Motohiro Yogo. Why do household portfolio shares rise in wealth? *The Review of Financial Studies*, 23(11):3929–3965, 2010. doi: 10.1093/rfs/hhq092.
- Cristobal Young, Charles Varner, Ithai Z. Lurie, and Richard Prisinzano. Millionaire migration and taxation of the elite: Evidence from administrative data. *American Sociological Review*, 81(3):421–446, 2016. doi: 10.1177/0003122416639625.
- 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.