# A levy cannot tax what its base cannot see

## The base caps the region, the rate moves you within it: redistribution as a parameter space

**Jason C. Braatz**
*Independent researcher*
jason@braatz.ai

**Preprint.** Posted 2026-08-29. Not peer reviewed, not submitted.

**Declaration of interest.** The author is employed by a company building accounting software for very small businesses. This work was conducted independently, on personal time, and without company funding, data or direction.

---

## Abstract

A multiplicative wealth process with a positive mean growth rate condenses: log-wealth variance grows without bound and the Gini approaches unity. **Every distribution that is not condensed is therefore being opposed by something.** This paper classifies the opposing mechanisms not by institutional origin but by four coordinates (**base, rate, periodicity, threshold**) plus the **realisation share** of the base, and asks which regions bound inequality below unity.

**The base sets a ceiling the rate cannot cross.** At a matched rate the two bases differ by roughly an order of magnitude in κ, the levy's compressive *budget*, for which the flow base admits a closed form, $\kappa = r \cdot E[\eta^+]$. But the budget is not the whole story, and the second half of it points the other way: **per unit of budget the flow base compresses more, not less.** Matched at $\kappa \approx 0.10$ the two reach Gini 0.222 and 0.125, and the reason is visible in a statistic nobody reports: the variance of the per-period log multiplier is untouched by the stock levy (a change of $6 \times 10^{-6}$) and falls by a third under the flow levy. **A stock levy truncates the outcome; a flow levy damps the generator**, and both register as a smaller Gini.

The stronger prediction this design was built to test (that a flow levy fails to oppose the multiplicative term *regardless of rate*) is **false**, and §3.1's rate sweep is what falsified it. The frontiers are **nested**, stock 0.000 against flow 0.125, not disjoint.

**The surviving claim is narrower and better: the decisive quantity is realisation** — the share of a period's gain the base can observe. At zero realisation a **100% levy on flow leaves the wealth vector exactly unchanged**, agent by agent, because its base is the uniform wage and a uniform assessment with a uniform rebate is the identity. And the collapse from there is violently front-loaded: **the first five per cent of the realisation axis covers 32% of the reachable range, and the first quarter covers 69%.** A rate is an intensity; realisation is an observability, and the observability binds first.

Periodicity and threshold are trim rather than structure: they modulate the effective rate without opening or closing a region, and a threshold at a quarter of the mean is close to free. A methodological result of independent interest: **a summary statistic with a hard ceiling cannot serve as a convergence criterion**: the Gini is capped at (N−1)/N, so a fully condensed economy also stops rising, and a drift test scores total condensation as bounded.

The claims are properties of a model class; no causal claim about any institution is made. All results reproduce from open code; 18 tests pin them.

**Keywords:** kinetic exchange models · wealth condensation · multiplicative growth · Gini coefficient · tax base · realisation · agent-based models · econophysics

**JEL classification:** D31, D63, H23, H24, C63

---

## 1 · Introduction

A multiplicative wealth process with a positive mean growth rate condenses. This is not a pathology of any particular model; it is what multiplicative processes do. The variance of log-wealth grows without bound, an additive wage becomes negligible in comparison, and the wealth share of the top holder tends to one. The kinetic-exchange literature has established this repeatedly and by several routes (Drăgulescu and Yakovenko, 2000; Patriarca, Chakraborti and Kaski, 2004; Yakovenko and Rosser, 2009).

The observation that motivates this paper is the contrapositive, and it is not often stated plainly:

> **Every wealth distribution that is not condensed is being opposed by something.**

Redistribution is not, in this framing, a policy question that arrives after the economics. It is a term in the process, and its *absence* is the special case.

That reframing makes a question available which the usual framing does not. Redistributive mechanisms are ordinarily compared by institutional identity (a progressive income tax, a wealth tax, zakat, a land value tax) and any comparison drawn on that axis is immediately and correctly read as an argument about which tradition is preferable. **But institutional identity is not what the process responds to.** The process responds to five numbers (the levy's four coordinates and the realisation share of its base) and it cannot see where any of them came from.

**The headline result is a ceiling, and it is not the one that gets argued about.** The base a levy is assessed on caps the region an economy can reach; the *rate* (the coordinate that receives essentially all public attention) moves an economy only within the ceiling its base has already fixed. A confiscatory levy on flow has approximately the compressive budget of a ten per cent levy on stock, because a levy cannot rebate a loss and so the gross positive growth rate caps what a flow base can assess at all.

**But the budget is not the mechanism, and this paper's sharpest finding is the gap between them.** Per unit of budget the flow base compresses *more* than the stock base, not less. §3.1 matches the two at $\kappa \approx 0.10$ and finds Gini 0.222 against 0.125, and then asks what each did to the *generator* rather than the outcome. The variance of the per-period log multiplier is 0.076542 unlevied, **0.076536** under the stock levy (a change of six parts in a million, which is to say none) and **0.051189** under the flow levy. A levy on stock rescales what a holder has and leaves the process that got them there exactly as it found it. A levy on flow reaches into the multiplicative term itself. **Both register as a smaller Gini, which is why the distinction is invisible in the statistic normally reported.**

**A prediction failed on the way here, and the failure is in the body rather than an appendix.** The claim this design was built to test was that a flow levy does not oppose the multiplicative term regardless of rate. That is false as stated, and the rate sweep is what falsified it: at full realisation a flow levy does bound the Gini, it is merely weak. The frontiers are **nested** (stock 0.000, flow 0.125), not disjoint.

**What survives is narrower, and it is the better result.** The quantity doing the work all along was **realisation**: the share of a period's gain the base is able to recognise. At $\rho = 0$ (the holder whose gains accrue and are never realised), a hundred per cent levy on flow leaves the wealth vector *exactly* unchanged, agent by agent. Not approximately, not to within a summary statistic: identically. The levy still fires; its base is simply the wage, the wage is the same for everyone, and a uniform assessment with a uniform per-capita rebate is the identity map.

That is the true *regardless of rate* result, and note what kind of statement it is. **It is a claim about what a base is able to observe, not about how hard it squeezes.**

**Contributions.**

1. A **parameterisation** of any period-assessed levy by four structural coordinates (base, rate, periodicity, threshold) plus the realisation share of the base, and a demonstration that the process's behaviour is a function of these alone (§2).
2. The result that **the base caps the reachable region and the rate only moves you within it**, with $\kappa$ (the levy's compressive *budget*, not its mechanism) separating the bases by an order of magnitude, and a closed form for the flow base's $\kappa$ that the simulation reproduces to within 7% at every rate tabulated (§3.1). The stock-versus-flow contrast this sharpens is prior and is credited in §3.1 and §6; what is new is $\kappa$ itself (the budget, separated from the mechanism) and the closed form.
3. **The outcome/generator distinction, with a witness for each** (§3.1.1). Two levies matched at the same budget compress unequally, and $\mathrm{Var}[\log a]$ says why: the stock base leaves the process that produces next period's distribution exactly as it found it. **An outcome measure records that the distribution was compressed; it does not record whether the mechanism producing the next one was touched.**
4. The identification of **realisation as the decisive quantity**, including the limiting result that a confiscatory levy on flow, at zero realisation, leaves the wealth vector exactly unchanged — its base is uniform, not absent (§3.2) — and the shape of the approach to that limit, which is steep enough that the first twentieth of the axis carries a third of the range.
5. A methodological result of independent interest: **a summary statistic with a hard ceiling cannot serve as a convergence criterion** (§3.4).
6. A **reproducible artefact**: every number below is regenerated from a public repository by the two commands §7 names (save the six quantities §7 enumerates, which no command prints) and the claims are held in place by the 18 tests in `tests/test_redistribution.py`, one of which exists specifically to make overclaiming fail loudly.

---

## 2 · The model

### 2.1 · The process

$N$ agents. Each period, every agent's wealth is multiplied by an idiosyncratic growth factor and receives a common additive wage:

$$\mathbf{w\_i(t{+}1) = w\_i(t){\cdot}(1 + \eta\_i(t)) + a}, \ \eta\_i \sim \mathcal{N}(\mu, \sigma^2)$$

The multiplicative term is the engine of condensation; the additive wage $a$ is the only force opposing it in the absence of a levy, and it is not enough. Throughout: $N = 800$, $\mu = 0.05$, $\sigma = 0.20$, $a = 0.05$, $T = 1200$ periods, with reported statistics averaged over the final quarter of the path.

### 2.2 · The levy, as four numbers

A period-assessed levy is described by:

| coordinate | meaning |
|---|---|
| **base** | what is assessed — the **stock** held, or the **flow** received |
| **rate** $r$ | the fraction of the liable amount taken |
| **periodicity** $P$ | periods elapsed between assessments |
| **threshold** $\theta$ | the exempt amount, in multiples of the mean of the base |

Everything collected in a period is redistributed per capita in the same period. The levy is therefore a **pure transfer**: aggregate wealth is untouched and only dispersion changes. This is verified to machine precision rather than assumed: `test_the_levy_is_a_pure_transfer` in `tests/test_redistribution.py` holds the implementation's reported `transfer_error` below 1e-12, so that no result below can be an artefact of the levy quietly changing the growth rate.

### 2.3 · Realisation

$\boldsymbol{\rho}$ is the share of a period's capital gain that is recognised as flow, and therefore enters a flow levy's base. $\rho = 1$ is a mark-to-market levy on accruals; $\rho = 0$ is the pure rentier whose gains accrue and are never realised.

$\rho$ is not a free parameter fitted to the result. It is a **stated structural property of every real tax system** (realisation-based taxation is not a modelling convenience but the near-universal practice) and it is **swept rather than chosen**: §3.1's flow rows are stated at $\rho = 1$ and labelled as such, §3.2 sweeps the axis, and the paper's central result is a statement about the whole $\rho$ axis rather than about a value of $\rho$ chosen to make it come out right.

### 2.4 · What is measured

**Gini**, of the wealth vector, in the exact sorted-rank form rather than a Lorenz approximation. **Top decile share**, because the Gini saturates (§3.4). And **$\boldsymbol{\kappa}$**, the share of aggregate wealth actually moved per assessment: the levy's *compressive budget.* It is a budget and not a mechanism: §3.1 matches two levies at $\kappa$ and finds them compressing unequally, and §3.3 removes a quarter of $\kappa$ at no measurable cost.

A fourth quantity is reported wherever the question is whether a levy reached the *generator* of the process rather than its outcome: **Var[log $\boldsymbol{a}$]**, the variance of the log per-period multiplier normalised by aggregate growth: written $a(\eta)$, a different object from §2.1's wage $a$, with which it unhappily shares a letter. It is a property of the process that produces next period's distribution, where the three statistics above are properties of this period's.

**Bounded** is the verdict the tables below carry, and it is a criterion rather than a statistic: a run is bounded when its Gini has settled over the final quarter of the path **and** its top decile sits below 0.90. The second condition is the one that does the separating, and §3.4 is the result that establishes why a settled Gini alone will not.

---

## 3 · Results

### 3.1 · The base sets a ceiling; the rate moves you within it

| levy | Gini | κ | top 10 % | bounded |
|---|---|---|---|---|
| none | 0.994 | — | 1.000 | **no** |
| stock, $r$ = 0.025 | 0.443 | 0.0250 | 0.336 | yes |
| stock, $r$ = 0.100 | 0.222 | 0.1000 | 0.193 | yes |
| flow, $r$ = 0.025 | 0.812 | 0.0025 | 0.734 | yes |
| flow, $r$ = 0.100 | 0.596 | 0.0102 | 0.481 | yes |
| flow, $r$ = 1.000 | 0.125 | 0.1026 | 0.138 | yes |

*The flow rows are assessed at full realisation, $\rho = 1$: §2.3's mark-to-market case and the implementation's default; §3.2 is the sweep that lowers it. These six rows are a selection: the rate sweep behind them is wider on both bases, and §3.4 quantifies over all of it.*

At a matched rate the two bases sit roughly an order of magnitude apart in $\kappa$, at every rate tested. The budget is the table's $\kappa$ column and is not a fitted relationship:

- for a **stock** base at zero exemption, $\kappa = r$ exactly: §3.3 raises the threshold and $\kappa$ falls;
- for a **flow** base, $\kappa = r \cdot \mathrm{E}[\eta^+]$, where $\mathrm{E}[\eta^+]$ is the *gross positive* growth rate, because a levy cannot rebate a loss. In closed form, $\mathbf{E[\eta^+] = \mu\Phi(\mu/\sigma) + \sigma\varphi(\mu/\sigma) = 0.1073}$ for the parameters above. The simulated $\kappa$ runs **4–7 % below** that form across the full rate sweep behind the table: −4.3 %, −4.6 %, −5.7 % at the three flow rates tabulated here ($r$ = 1.000, 0.100, 0.025), reaching −6.8 % at the sweep's lowest rate, $r$ = 0.010. The residual is flat between $r$ = 1.000 and $r$ = 0.500 and widens monotonically below it, which makes it a denominator convention rather than noise: the implementation measures $\kappa$ against post-growth wealth. *These residuals are computed from the unrounded $\kappa$ rather than from the four-decimal values the table displays; at* $r^*$ = 0.025 that display quantum is ±2 % of $\kappa$ itself, which is wider than the spread being reported.* The test suite asserts agreement within 10 %.

So a **confiscatory** levy on flow has approximately the compressive budget of a **10 % levy on stock**. The base is not a detail of implementation. It sets a ceiling, and the rate (the coordinate that receives essentially all public attention) moves an economy only within the ceiling its base has already fixed.

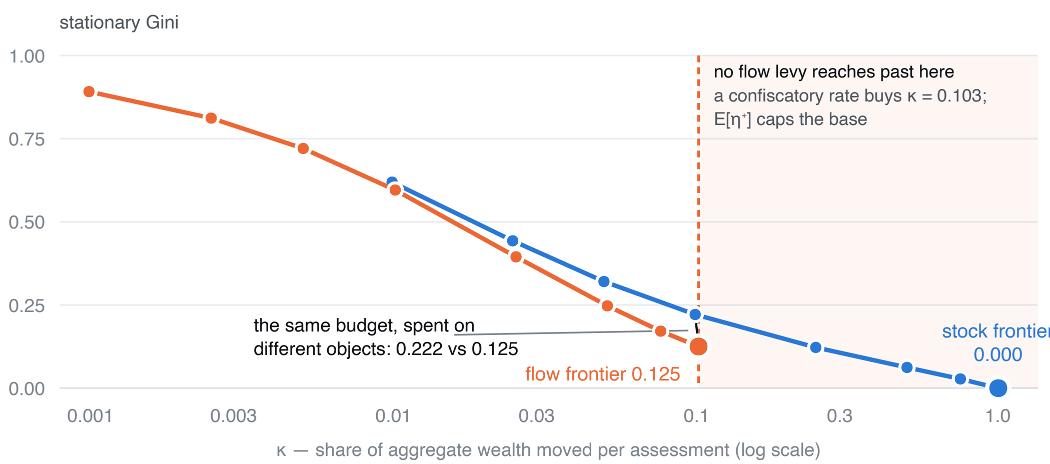


**Figure 1 — Nested frontiers, not disjoint ones.**

Stationary Gini against κ, both bases, log budget axis. Two facts sit in one picture: the flow curve lies *below* the stock curve everywhere the two overlap (more compression per unit of budget) and the flow curve *stops* at κ = 0.103, because $E[\eta^+]$ caps what the base can assess. The frontiers are nested, stock 0.000 inside flow 0.125, and the nesting is what falsified the stronger prediction below.

### 3.1.1 · The two bases act on different objects, and only one statistic can tell

**The two bases do not merely differ in budget. They act on different objects.** Matched at κ ≈ 0.10, the two levies compress the cross-section unequally (Gini 0.222 against 0.125) but the more telling comparison is what each does to §2.4's fourth quantity, Var[log *a*]: the generator of the process rather than its outcome. Unlevied, Var[log *a*] = 0.076542. Under the **stock** levy at that budget it is **0.076536**: a change of $6 \times 10^{-6}$, which is to say none at all. Under the **flow** levy it is **0.051189**, a third lower. A levy on stock rescales what a holder has and leaves the process that got them there exactly as it found it; a levy on flow reaches into the multiplicative term itself. The stock base truncates the outcome, the flow base damps the generator, and both register as a smaller Gini, which is why the distinction is invisible in the statistic normally reported. An outcome measure records that the distribution was compressed. It does not record whether the mechanism producing next period's distribution was touched.

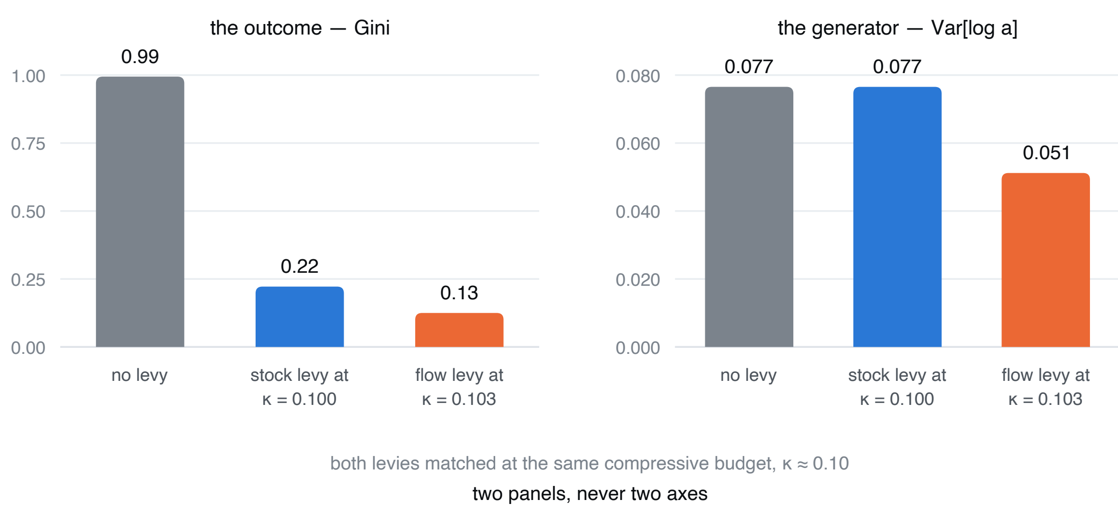


**Figure 3 — One statistic falls for both. The other**

**does not.** Two panels: never two axes, since the quantities carry different units and different scales. Left: the Gini, which falls under both levies. Right: Var[log $a$], which is flat under the stock levy and a third lower under the flow levy. The picture is the argument: a reader looking only at the left panel would call the two levies comparable at this budget.

**This contrast is not new.** Bouchaud and Mézard (2000) carry a flow levy, a stock levy and the per-capita redistribution of each in a single wealth balance, and give the stationary Pareto exponent in closed form in all four of their own tax parameters: a rate and a redistributed fraction for each base. They write that exponent μ: a different object from §2.1's growth drift μ, with which it unhappily shares a letter. Their ranking is the one measured here, and they state it more strongly: income taxes *"tend to reduce the inequalities of wealth (i.e., lead to an increase of μ), even more so if part of this tax is redistributed"*, while *"quite surprisingly, capital tax, if used simultaneously to income tax and not redistributed, leads to a decrease of μ"*. Their stock levy can *reverse* the sign of the effect; the one measured here merely buys less compression per unit of budget. What this section adds is not the contrast but the pair of witnesses for it (κ, which says how much budget a base has, and Var[log $a$], which says whether the levy spent it on the outcome or on the generator) in a discrete process where the two can be matched and separated. §6 states what that leaves.

**A prediction that half-failed.** The claim this section was built to test was stronger: that a levy on flow does not oppose the multiplicative term *regardless of rate.* That is **false as stated**, and the sweep is what falsified it. At full mark-to-market realisation a flow levy does bound the Gini; it is merely weak. Rate 1.00 on flow reaches Gini 0.125, which a stock levy reaches at rate 0.25. The reachable frontiers are **stock 0.000 < flow 0.125**: the bases occupy *nested* regions, not disjoint ones. The surviving claim is narrower and is in the next section.

## 3.2 · Realisation is the crux

The multiplicative term operates on the **stock**. A flow base reaches it only through whatever share of the period's gain is recognised as income. That share is ρ, and it is the quantity that was doing the work all along:

| realisation ρ | reachable Gini (flow base) | share of the reachable range covered |
|---|---|---|
| 0.00 | **0.994** | 0.0% |
| 0.05 | 0.720 | **31.5%** |
| 0.10 | 0.596 | 45.9% |
| 0.15 | 0.509 | 55.8% |
| 0.25 | 0.395 | **69.0%** |
| 0.40 | 0.293 | 80.7% |
| 0.55 | 0.229 | 88.1% |
| 0.70 | 0.184 | 93.3% |
| 0.85 | 0.150 | 97.1% |
| 1.00 | 0.125 | 100.0% |

*The 0.00, 0.25 and 1.00 rows are the three this paper previously reported; the seven between them are new to this draft, run on the same committed model at the same seed and the same horizon. They were added to draw Figure 2 and they changed what §3.2 claims, which is why they are tabulated rather than only plotted.*

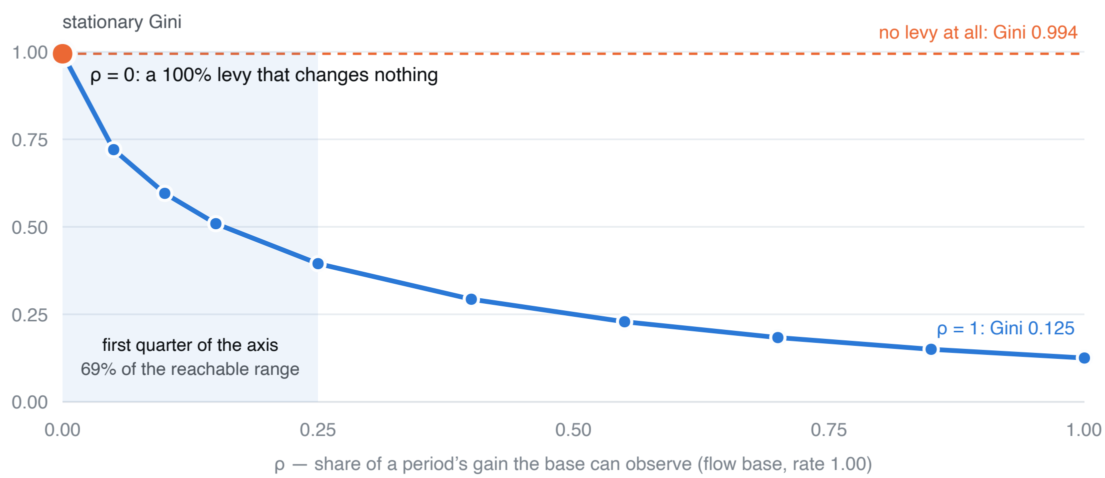


**Figure 2 — Realisation is the crux, and its first five per cent does a third of the work.** The reachable Gini against ρ under a confiscatory flow levy, with the no-levy Gini drawn as a horizontal reference that the ρ = 0 point sits exactly on.

**The shape of that column is a second result and it was not visible from three rows.** The collapse is violently front-loaded: **the first twentieth of the realisation axis covers 31.5% of the whole reachable range, and the first quarter covers 69%.** Going from a base that sees nothing to a base that sees five per cent of a period's gain buys more compression than going from a quarter to everything.

That has a reading and the reading should be stated carefully, because it is the kind of sentence this paper's §5 exists to restrain. **Inside this model class**, the marginal return to widening what a base can observe is enormously higher at the closed end of the axis than at the open end. Whether any real assessment regime sits at the closed end, and what it would cost to move it, are questions this paper does not touch and has no evidence about. What it can say is that the axis is not linear, that the existing three-row reading understated how much is concentrated at the bottom of it, and that a design comparing realisation regimes should not assume equal steps buy equal compression.

At $\rho = 0$ (the holder whose gains accrue but are never realised), a **100 % levy on flow leaves the wealth vector exactly unchanged**: Gini 0.994 against 0.994, top decile 1.000 in both, and the two paths agree agent by agent rather than merely on the summary statistics. The identity is structural. The levy is still assessed (at $\rho = 0$ the flow base is not empty but is the accrued **wage**, and the assessments do fire) but the wage is identical for every agent, so the levy takes the same amount from each and returns it per capita. A uniform assessment with a uniform rebate is the identity on the wealth vector. What $\rho = 0$ removes is not the levy but the **dispersion in its base**.

That is the true "regardless of rate" result, and note what kind of statement it is. It is a claim about what a base is able to **observe**, not about how hard it squeezes. A rate is an intensity; realisation is an *observability*, and the observability binds first.

Unrealised appreciation is wealth whose growth the assessing layer has not been asked to recognise, and what the $\rho$ axis measures is how much of a period's gain that layer can see at all.

**A stronger reading was available, was tested, and is withdrawn here.** One might hold that a levy which cannot see an accrual and a financial statement which does not record a degradation *are the same structure*, seen from two sides. Put to a cross-scale check against the companion work on the reporting layer, that identification does not hold. The share of a change a reporting filter fails to recognise is **deferred** (it accumulates and is released later at a stated rate) while the share of a gain this model's base fails to recognise is **never assessed at all**. Deferred arrival and non-arrival are different operators, and a shared adjective between them is not an equation. The check, its thresholds, and its preregistration (the withdrawal was registered before the check was run) are recorded in `docs/RESULT-END-TO-END-001-E1.md`.

### 3.3 · Periodicity and threshold are trim, not structure

Both remaining coordinates modulate the *effective rate* rather than opening or closing a region, which is what leaves base and rate as the two structural axes.

**Periodicity.** On a stock base, holding the average rate constant at 0.02 per period, assessing every $P$ periods at rate $0.02{\cdot}P$ moves the stationary Gini from 0.486 ($P = 1$) to 0.456 ($P = 20$). A lumpier assessment is very slightly **stronger** over that range, because it catches dispersion that has had time to accumulate. The effect is not monotone in $P$, and the sweep behind those two endpoints says so: the minimum is **interior**, 0.451 at $P = 30$, and by $P = 50$ (where holding the average rate at 0.02 requires the maximum rate, 1.00) the Gini has returned to 0.469, above its $P = 20$ value. The whole sweep spans 0.035, which is the operative fact: an annual assessment is not a watered-down continuous one.

**Threshold.** On the same base at $r = 0.025$, monotone and smooth, with no cliff: Gini 0.443 at zero exemption rising to 0.770 at 20× the mean of the base. The interesting part is the near end. A threshold at **0.25× the mean costs nothing measurable** in compression (0.444 against 0.443) while reducing κ by a quarter. Exempting small holders removes a quarter of the assessed volume and none of the effect, because the compression is performed by transfers at the top of the distribution.

A threshold that exempts the poor is therefore not a concession that weakens the mechanism. It is close to free.

**κ is necessary and it is not sufficient, and there is a witness on each side.** §3.1 matches the two bases at $\kappa \approx 0.10$ and finds them compressing unequally, 0.222 against 0.125; the threshold sweep above supplies the converse from the other side, removing a quarter of κ at no measurable cost in compression, 0.444 against 0.443. κ can hold while the outcome moves and move while the outcome holds, so **no function of κ alone reproduces §3.1's table.** κ is what a base makes available to spend (which is why the bases sort, and why the closed form is worth having) but what the spending buys is fixed by the object the levy acts on, which is the distinction §3.1's two-witness comparison draws.

### 3.4 · A methodological result: saturating statistics cannot detect convergence

The boundedness criterion was first written as a drift test: the Gini has settled if its mean over the last quarter of the path exceeds the previous quarter's by less than a tolerance. It scored the **unopposed** process, the one the entire exercise exists to contrast against, as *bounded.*

The Gini of $N$ agents is capped at $(N-1)/N$. A fully condensed economy therefore also stops rising, not because it reached a stationary distribution but because it ran out of headroom. At $N = 800$ and $T = 1200$ the unopposed process reads Gini 0.994 and flat (short of the 0.99875 ceiling it is pinned against) while its top decile holds 1.000 of everything. The drift test was measuring the ceiling.

§2.4's criterion pairs the settled Gini with a top decile below 0.90 for exactly this reason, and it is the second condition that does all of the separating. Across §3.1's full rate sweep (wider than the six rows tabulated there), the bounded runs' Gini spans 0.000–0.891 against the condensed run's 0.994: a gap of 0.103 whose upper edge is a *saturated reading* and not the 0.99875 ceiling it falls short of, so any Gini threshold would have to be drawn inside it and redrawn for every $N$; their top decile spans 0.100–0.861 against 1.000, clearing the 0.90 threshold with 0.039 to spare.

**The general rule, which is not confined to this model: a summary statistic with a hard ceiling cannot distinguish "converged" from "saturated."** Before using one as a convergence criterion, ask what its maximum is and whether the failure mode you are trying to detect drives it there.

---

## 4 · Abandoned approaches

*A result reported without the routes that failed is a result the reader cannot calibrate, they are shown the one path that worked and left to assume it was the only one considered. Everything here was actually attempted.*

**Classification by institutional origin.** The first framing compared named systems. It was abandoned because it is unanswerable as posed: any ranking of institutions is read as an argument about traditions, and correctly so. The parameter-space framing replaces "which system do you favour" with "which regions bound the Gini below unity", which is a question the models answer and no one can call advocacy. The cost is that the paper can no longer say anything about any actual institution, and it does not.

**"Regardless of rate."** §3.1. The original claim was sharper and false. The shape of the failure is informative: the direction was right, the mechanism was misidentified, and the surviving claim is the narrower one about realisation.

**Defining "flow" so that the claim came out right.** Once "regardless of rate" failed, an obvious repair was available: redefine the flow base so that it excludes the components that rescued it. This was refused. A definition adjusted until the prediction survives is a free parameter wearing different clothing. The claim was narrowed instead.

**The drift-only boundedness test.** §3.4. It looks like a convergence check. It was caught by asking what the statistic's maximum was, and it is now pinned by a test named `test_a_flat_gini_does_not_mean_a_bounded_one`, so that any future simplification of the criterion fails loudly instead of quietly re-scoring condensation as success.

---

## 5 · Limitations

1. **ρ is exogenous here, and in the world it is not.** Realisation is chosen, and it responds to the rate: raising a flow levy gives holders a reason to realise less. Endogenising ρ would make the flow base *weaker* than reported, but it is unmodelled, and a reader should treat the ρ axis as a comparative static rather than a policy response function.
2. **This is a model-class result.** No causal claim about any historical or contemporary institution is made, and none is supported. No field evidence is used, required, or available.
3. **No production, no labour supply, no portfolio choice, no behavioural response to the levy.** The Lucas critique applies in full force and is not answered here.
4. **One good, one asset, no prices.** The process is a wealth process, not an economy.
5. **Finite N, one seed per reported figure, and a fixed parameter neighbourhood.** $N = 800$, and every *simulated* number above is a mean over a tail window of a **single** path at `seed = 0` rather than an ensemble average: the exceptions are the five closed-form quantities §7 names: §3.1's $E[\eta^+]$ and its three $\mathrm{Var}[\log a]$ values, which are quadrature, and §3.4's Gini ceiling, which is arithmetic in $N$. Seed-robustness is asserted separately rather than averaged in: `test_the_result_is_not_a_lucky_seed` holds two configurations inside their stated bands across five seeds, at the reported $T = 1200$ as well as at the suite's $T = 600$. The qualitative separations are large relative to those bands, but the third decimal is not defended.
6. **The realisation sweep's shape rests on one seed, like everything else here.** §3.2's seven new rows are a single path at `seed = 0` at $T = 1200$, the paper's stated convention.

The front-loading is large relative to the seed bands `test_the_result_is_not_a_lucky_seed` holds (the $\rho = 0.05$ row sits 0.27 below the $\rho = 0$ row) but the exact percentages of the range are not defended to the third decimal, and a reader wanting the shape rather than the numbers should read Figure 2 rather than the table.

7. **Growth shocks are Gaussian and i.i.d. across agents and time.** Heavy tails and autocorrelation both plausibly matter and neither is explored.
8. **Positive, not normative — and that boundary is a constraint on the reader too.** Nothing here implies that bounding inequality is desirable. It implies only that certain regions of the parameter space do it and others do not.

---

## 6 · Relation to existing work

The condensation result is standard in kinetic exchange (Chakrabarti, Chatterjee, Chakravarty and the surrounding literature), where the effect of saving propensity, taxation and redistribution on stationary wealth distributions has been examined from several directions. **Two works in that literature are prior to this paper's central contrast.**

**Bouchaud and Mézard (2000)**, whose wealth balance and closed-form Pareto exponent §3.1 sets out, are prior to the stock-versus-flow ranking used there. The contrast between the two bases (in terms of what each does to the shape of the stationary distribution) is theirs, and the per-capita rebate fraction is a parameter in their solution rather than an extension awaiting one. Their solution is continuous-time and carries neither a periodicity nor a threshold, so §3.3's two trim coordinates are outside it.

**Benhabib, Bisin and Zhu (2011)** supply three further results that bound what is left. Their Proposition 3 has the tail index rising in both the estate tax and the capital income tax, so the *nested* frontiers this paper reaches in §3.1 were already visible in a different metric and a different model. Their Proposition 4 has tail inequality rising with a mean-preserving spread of the return process, which is the general form of §3.1's finding that the flow levy reaches the dispersion of the multiplier and the stock levy does not. And their §4.1 notes that an economy whose multiplier is bounded below one has a stationary distribution bounded above, with no power-law tail at all: a claim this paper does not make and does not need, but the first one any extension of §3.1 toward tail indices, the object Gabaix (2009) surveys, would meet.

**What remains is narrower than the contrast.** It is not that redistribution opposes condensation, and it is not that the two bases act differently on the shape of the distribution. It is that the mechanisms sort by **observability of the base** (§3.2) rather than by rate or institutional form; that the budget through which they operate has a closed form ($\kappa$) rather than being a simulation regularity, though the sorting is not a function of that budget alone (§3.1); and that a single sweep separates the two by measuring the generator and the outcome side by side. Three further differences are of construction rather than of claim, and none is offered as a result: the levy here is on the **realised gain only**, with no loss offset, so the multiplier is asymmetrically truncated rather than symmetrically contracted toward one; the two bases are compared at matched compressive **budget** rather than at matched rate; and the process is a discrete Kesten-type recursion with an explicit per-period budget identity rather than a continuous-time mean-field one.

The realisation result touches the public-finance literature on the choice of base (Kaldor, 1955) and on realisation-based versus mark-to-market taxation (Auerbach, 1991; Toder and Viard, 2016; Saez and Zucman, 2019) from an unfamiliar angle: not from the incentive or valuation side, but as a statement about the information available to the assessing layer. On the stock-versus-flow axis, the paper is silent about optimal taxation and deliberately so; it characterises reachable regions, not desirable ones.

---

## 7 · Data and code availability

**The seven new realisation rows in §3.2, and Figures 1–3, are produced by `scripts/wt221_paperII_figure_data.py`**, which imports the same committed model this paper has always used and writes the plotted values to JSON. It is committed alongside this draft, and it takes no arguments: the seed, the horizon and the parameter neighbourhood are those §2.1 states.

All results in this paper are produced by open code and no proprietary or restricted data is used, because no empirical data is used at all: every measured number is generated by simulation, save the four closed-form quantities the next bullet names and §3.4's Gini ceiling $(N-1)/N = 0.99875$, which is arithmetic in $N$ and is printed by no command here.

- **Repository:** `https://github.com/jasoncbraatz/wealth-tensor` (public)

- **Module:** `src/wealth_tensor/redistribution.py`
- **Regenerate every number in §3:** `python3 scripts/wt030_report.py`: except §3.1's four closed-form quantities: $E[\eta^+] = 0.1073$ and the three Var[log $a$] values, which are quadrature over the multiplier's distribution rather than simulation output and come from `python3 scripts/wt077_tail_index.py`, and except six quantities neither command prints in any precision: §3.4's Gini ceiling $(N-1)/N = 0.99875$, which is arithmetic in $N$; §3.4's 0.90 top-decile criterion, which is a chosen threshold and not an output; §3.3's 0.035 periodicity span and §3.4's 0.103 Gini gap, each a difference of two values `wt030_report.py` prints; §3.4's 0.039 top-decile margin, the distance from that command's printed 0.861 to the 0.90 threshold above; and §3.1's $6 \times 10^{-6}$ change in Var[log $a$], the difference of two values `wt077_tail_index.py` prints.
- **Test suite:** `python3 -m pytest tests/ -q` runs the whole repository; the **18** tests in `tests/test_redistribution.py` are the ones that hold this paper's claims in place, and that count is the one quoted in the abstract and in §1.
- **The two tests that make overclaiming fail loudly here — the suite holds others:** `test_a_flat_gini_does_not_mean_a_bounded_one`, which pins §3.4's boundedness criterion so that any future simplification of it fails instead of quietly re-scoring condensation as success, and `test_excess_demand_is_monotone_here_so_this_is_not_an_SMD_result`, which constrains this programme's price-formation manuscript (that manuscript is superseded and withdrawn, and the guard outlives it) in a companion module of the same suite, `tests/test_excess_demand.py`; it is a different companion from §3.2's work on the reporting layer.
- **Commit for the results reported here: 3b11f23**: the last commit touching `src/wealth_tensor/redistribution.py`, and therefore the state of the module that produced §3's simulation output. The pin is **per file** deliberately: a pin on the last commit touching `src/` as a whole is a sentence whose truth changes whenever any unrelated module moves, and nothing in the repository watches it. It does **not** cover the two `scripts/` commands named above, which produce §3 numbers from outside `src/`. *A head-of-repository SHA will additionally be pinned when this paper is posted, and it is what covers them.*

**Known limitations of this review.** Every number above is regenerated by the two commands named here, and a test suite pins the claims; the cross-references, the citation pointers and the reference entries were checked by a sequence of independent reads rather than mechanically, and the simulation figures rest on a single seed per configuration with seed-robustness asserted separately rather than averaged in. Two things were not checked at all and are named rather than left to be found: §1's characterisation of zakat is not supported here by a primary source, and §3.1's "*4–7 %*" band is an under-claim: the sweep's largest residual is 6.831 %, so the band is true as written and looser than the evidence requires.

Pinning the last commit that touched the module rather than a bare placeholder is deliberate: it is non-circular (a paper cannot cite the commit that adds the paper), it is verifiable today, and it names the object a replicator actually needs: the state of the code, not the state of the prose.

The repository's `docs/` directory is deliberately public and contains the project's working notebook, including the entries in which the claim of §3.1 half-failed and was narrowed. It is part of the record rather than an appendix to it.

---

## References

*Bibliographic details for the entries marked ✓ were verified against live sources on 2026-08-10. The two marked ✓⧗ were re-verified against their Crossref records on 2026-08-17 and name, in the entry, the pre-publication version actually read, per* `REFERENCE-POLICY` *§4. The remainder are standard works whose details are to be re-checked at submission.*

**Kinetic exchange and wealth condensation**

Bouchaud, J.-P., & Mézard, M. (2000). Wealth condensation in a simple model of economy. *Physica A: Statistical Mechanics and its Applications*, 282(3–4), 536–545. `doi:10.1016/S0378-4371(00)00205-3` ✓⧗ *(issue and pagination checked against the Crossref record, 2026-08-17. Text consulted: arXiv* `cond-mat/0002374`*, read in full. The quotations in §3.1 are attributed to that preprint and may not appear verbatim in the article of record. Consulted 2026-08-17 / published 2000.)*

Chakrabarti, B. K., Chakraborti, A., Chakravarty, S. R., & Chatterjee, A. (2013). *Econophysics of Income and Wealth Distributions.* Cambridge University Press.

Chakraborti, A., & Chakrabarti, B. K. (2000). Statistical mechanics of money: how saving propensity affects its distribution.

*The European Physical Journal B*, 17(1), 167–170. ✓

Chatterjee, A., & Chakrabarti, B. K. (2007). Kinetic exchange models for income and wealth distributions. *The European Physical Journal B*, 60(2), 135–149.

Drăgulescu, A., & Yakovenko, V. M. (2000). Statistical mechanics of money. *The European Physical Journal B*, 17(4), 723–729. ✓

Patriarca, M., Chakraborti, A., & Kaski, K. (2004). Statistical model with a standard Γ distribution. *Physical Review E*, 70, 016104.

Yakovenko, V. M., & Rosser, J. B. (2009). Colloquium: Statistical mechanics of money, wealth, and income. *Reviews of Modern Physics*, 81(4), 1703–1725. ✓

**Wealth dynamics and inequality**

Benhabib, J., Bisin, A., & Zhu, S. (2011). The distribution of wealth and fiscal policy in economies with finitely lived agents. *Econometrica*, 79(1), 123–157. `doi:10.3982/ECTA8416` ✓⧗ *(page range checked against the Crossref record, 2026-08-17. Text consulted: NBER Working Paper 14730 full text, read in full; §6's characterisation of Propositions 3 and 4 and of §4.1 is taken from that version and the numbering may differ in the article of record. Consulted 2026-08-17 / published 2011.)*

Gabaix, X. (2009). Power laws in economics and finance. *Annual Review of Economics*, 1, 255–294.

Gini, C. (1912). *Variabilità e Mutabilità*. Tipografia di P. Cuppini.

**Public finance: base, realisation, and mark-to-market**

Auerbach, A. J. (1991). Retrospective capital gains taxation. *American Economic Review*, 81(1), 167–178.

Kaldor, N. (1955). *An Expenditure Tax*. George Allen & Unwin.

Saez, E., & Zucman, G. (2019). Progressive wealth taxation. *Brookings Papers on Economic Activity*, 2019(2), 437–533. ✓

Toder, E., & Viard, A. D. (2016). Replacing corporate tax revenues with a mark-to-market tax on shareholder income. *National Tax Journal*, 69(3), 701–731.

**Methodological**

Lucas, R. E. (1976). Econometric policy evaluation: a critique. *Carnegie-Rochester Conference Series on Public Policy*, 1, 19–46.

---

*§1's characterisation of zakat (a levy assessed on stock held above a threshold across a full year) is not supported here by a primary source. The paper's argument does not depend on it: the institution is a measurement, not a premise, and §1 says so.*

---

## Use of AI assistance

Anthropic Claude Opus 5, at high reasoning effort, was used throughout as a research and drafting assistant: literature retrieval, adversarial review, code review and prose drafting. All claims, results and final text are the author's, and every computational result is produced by committed code in the repository named in §7.

---