The Wealth of Winners:
AI Compute and the Compounding Lottery

Efraim Berkovich*

March 3, 2026

WORK IN PROGRESS

Abstract

When AI agents drive the frontier of innovation, capital accumulates through winner-take-all discovery shocks rather than traditional thrift. Analysis of a general-equilibrium "lottery economy" reveals a single technological parameter, γ–the elasticity of discovery returns with respect to compute spending–governs the morphology of long-run inequality. Sublinear elasticity ($\gamma < 1$) ensures an ergodic economy with bounded inequality; unitary elasticity ($\gamma = 1$) generates Pareto tails through random growth; and superlinear elasticity ($\gamma > 1$) triggers condensation and the emergence of a persistent oligarchic state. Endogenous compute spending creates a "preemption motive" that can overturn mean reversion: even risk-averse agents engage in discovery arms-races that consolidate dynastic wealth. In AI-driven economies, the technological return to scale in discovery, rather than preferences or entrepreneurial skill, becomes the primary determinant of the wealth distribution and social mobility.

Keywords: AI and economic growth; wealth inequality; random growth processes; innovation lotteries; general equilibrium; endogenous search; winner-take-all markets

JEL Codes: D31, O33, O41, D63

*Email: eqberkovich@gmail.com

1 Introduction

As artificial intelligence (AI) assumes a dominant role in production, the primary engine of wealth accumulation is shifting. In standard macroeconomic models, capital is accumulated through thrift: households save a portion of their wage and capital income to gradually build wealth. By contrast, in an AI-driven economy, the source of wealth creation is not the incremental accumulation of physical stock, but the discovery of high-leverage innovations–algorithms, molecules, or strategies–found by automated agents.

This AI economy is a risky search environment. Rather than merely renting capital, households allocate resources to AI agents to search for profitable innovations. This process is inherently stochastic and lumpy: while the majority of search attempts fail, a successful discovery yields a winner-take-all return. In a tractable general-equilibrium framework where capital accumulation is governed exclusively by this lottery mechanism, acentral question arises: distinct from preferences or talent, how does the technology of discovery shape the distribution of wealth when luck is a core driver of wealth accumulation?

When returns to innovation are large, the shape of the distribution is dominated by lottery effects, even as other, more standard, mechanisms of wealth accumulation are present. Relative wealth is reshuffled through a risky allocation mechanism. Inequality arises from who absorbs downside risk and who captures upside, not from thrift. Importantly, wealth formation in this economy is not meritocratic in the conventional sense: accumulation reflects access to AI compute rather than entrepreneurial effort or individual skill.

Analysis isolates a single sufficient statistic: the elasticity of access to innovation (γ) with respect to wealth. This parameter captures how effectively money translates into discovery probability. This technological primitive governs not only the tail behavior of wealth, but the very mechanism of accumulation: as γ varies, the economy transitions across distinct regimes characterized by fundamentally different incentives and competitive dynamics. This paper makes the following contributions:

1. **A unified framework for discovery-based growth.** A tractable general-equilibrium framework is developed that nests distinct processes–specifically additive jump processes and multiplicative random growth–within a single model governed by the elasticity parameter γ.

2. **Characterization of long-run inequality.** The existence and uniqueness of a stationary distribution for concave discovery functions are established, with analytical derivations of tail properties across regimes. It is shown that sublinear scaling ($\gamma < 1$) yields bounded inequality, linear scaling ($\gamma = 1$) recovers Pareto tails, and superlinear scaling ($\gamma > 1$) triggers wealth condensation.

3. **The Preemption Motive and arms-races.** Numerical characterization of the equilibrium reveals a series of phase transitions as γ increases. A key finding is the emergence of a "preemption motive" in the arms-race regime: the prospect of future dominance neutralizes standard precautionary savings behavior, inducing socially excessive investment in discovery even among risk-averse households.

4. **Welfare and policy benchmarks.** Decentralized outcomes are contrasted with the centralized pooling of discovery risk. Using Consumption Equivalent Variation metrics, the welfare costs of decentralized discovery are quantified, highlighting the efficiency losses inherent in non-cooperative innovation lotteries.

1.1 Related Literature

This paper situates the "lottery economy" at the intersection of macroeconomic literature on wealth inequality, innovation-driven growth, and the transformative role of artificial intelligence.

First, the analysis builds on the stochastic growth literature that characterizes heavy-tailed distributions through random growth processes (Gabaix, 1999; Luttmer, 2007). While these classic models typically rely on Gibrat's Law–where growth is proportional to size–to generate Pareto tails, Benhabib et al. (2011) demonstrate that idiosyncratic capital income risk is a primary driver of the wealth tail in a standard heterogeneous-agent environment. The framework presented here treats the elasticity of discovery access (γ) as the governing primitive. This approach nests additive and multiplicative growth as specific regimes: sublinear scaling restores mean reversion, while superlinear scaling produces non-ergodic condensation. By endogenizing the growth rate through compute investment, the model provides a structural basis for the transition between these distributional morphologies.

Second, the paper relates to Schumpeterian growth models connecting innovation to top wealth shares (Jones and Kim, 2018; Aghion et al., 2019). While that literature focuses on entrepreneurial effort or firm-level heterogeneity, this paper abstracts to a mapping from capital resources to discovery probabilities. The elasticity of this mapping is shown to be a sufficient statistic for whether innovation reallocates wealth in an ergodic manner or facilitates persistent dynastic dominance. This complements findings by Hubmer et al. (2023), who identify heterogeneous returns to capital as a primary driver of top wealth shares.

Third, the mechanism described here differs from standard models of entrepreneurship and occupational choice (Cagetti and De Nardi, 2006). In traditional Bewley-Aiyagari environments (Bewley, 1977; Aiyagari, 1994), wealth accumulation is driven by thrift and precautionary motives. In contrast, the AI-driven economy presented here operates through

a *preemption motive*, where wealth is deployed as an entry fee into winner-take-all innovation lotteries. While this aligns with empirical evidence regarding wealth-contingent access to high-return opportunities (Chodorow-Reich et al., 2024), the model formalizes this through a contest success function (Tullock, 1980) embedded in a dynamic general equilibrium. The key distinction from the entrepreneurship literature is that accumulation is driven not by effort or talent, but by a household's AI compute spending, a financing of luck in the innovation lottery.

Fourth, this work contributes to the growing literature on the macroeconomics of AI. While much of the existing research focuses on labor displacement and automation (Acemoglu and Restrepo, 2018), this paper follows the intuition of Korinek (2024) that AI may amplify scale effects, providing a formal general-equilibrium framework for this intuition. By modeling AI as a discovery technology, the paper demonstrates that when search exhibits increasing returns, inequality is governed primarily by technological access rather than by patience or risk attitudes. This model suggests that, in an AI-driven economy, the primitive governing long-run inequality is no longer the saving propensity (Piketty, 2014), but the technological scaling of the search process.

Finally, the results connect to the physics of condensation and preferential attachment (Krapivsky and Redner, 2001). By providing an economic microfoundation for superlinear reinforcement, the model shows how the "Wealth buys Compute" feedback loop leads to the emergence of dominant agents and captures the general equilibrium effects of such concentration on welfare and social mobility.

Overall, the paper reframes wealth inequality in innovation-driven economies as a problem of access to discovery rather than thrift or talent. Isolating the elasticity of discovery access as the central technological primitive provides a unifying perspective on random growth, entrepreneurship, and AI-driven concentration.

2 Model

In a dynamic economy model consisting of N infinitely lived households, time is assumed discrete. Aggregate output follows a Cobb–Douglas production function, $Y_t = A_t K_t^{\alpha} N^{1-\alpha}$, where K_t is the aggregate capital stock and households supply labor inelastically, and $A_t = 1$ is fixed. With α as the capital share, the rental rate of capital is $r_t = \alpha(K_t/N)^{\alpha-1}$ and the wage is $w_t = (1-\alpha)(K_t/N)^{\alpha}$.

The central feature of the model is the investment technology. Capital accumulation is mediated by a stochastic "innovation lottery" characterized by creative destruction. Interpret $\delta \in (0,1)$ as the rate of technological obsolescence: the arrival of new technology renders a fraction of existing capital less valuable.

In each period t, a single innovation event occurs that embodies the period's technological progress in a new capital vintage of value Π_t. Property rights to this innovation are awarded via a winner-take-all contest. Households allocate resources $m_{i,t}$ to purchase compute time for AI agents. AI agents search for innovation, with compute spending mapping to search intensity. One AI agent finds the innovation and thereby produces the innovation prize Π_t to the funding household. Thus, purchases of compute are effectively lottery tickets.

Under the assumption that the period prize to equal aggregate search outlays,

$$\Pi_t \;=\; \sum_{i=1}^{N} m_{i,t}. \tag{1}$$

Interpretation of this assumption is in Section 2.2. Under the *forced participation* benchmark analyzed in Section 2.3, $m_{i,t} = r_t k_{i,t}$ so that $\Pi_t = r_t K_t$. In the case of the endogenous choice economy in Section 2.4, $\Pi_t = \sum_i m_{i,t}$ in general differs from $r_t K_t$.

Let $L_t \in \{1, \ldots, N\}$ be a random variable denoting the index of the unique winning household at time t. The probability distribution of L_t is a categorical distribution governed by the relative search intensities of the households:

$$\mathbb{P}(L_t = i) = p_{i,t} = \frac{\pi(m_{i,t})}{\sum_{j=1}^{N} \pi(m_{j,t})} \tag{2}$$

where $\pi(\cdot)$ represents the impact function (or contest success function). By construction, $\sum_i p_{i,t} = 1$, ensuring exactly one winner per period.

The law of motion for the capital stock of household i is:

$$k_{i,t+1} = (1 - \delta)k_{i,t} + \mathbb{I}_{\{L_t = i\}}\Pi_t \tag{3}$$

where $\mathbb{I}_{\{\cdot\}}$ is the indicator function. Under this specification, all households suffer a relative wealth loss due to obsolescence, while the winning household receives a discrete jump in their capital share by capturing the new vintage.

Remark 1. *Aggregate Growth A model expanstion can assume A_t grows deterministically. As shown in Appendix A, normalizing capital by the growth factor $A_t^{1/(1-\alpha)}$ reduces the dynamics to a stationary aggregate baseline. Analysis proceeds with these normalized variables, effectively setting $A_t \equiv 1$ to focus on distributional dynamics.*

Interpretation of the impact function π. Parameterize the effective lottery weight as $\pi(m) = m^\gamma$, where $\gamma \geq 0$ is the *elasticity of access* and depends on compute spending m.

- If $\gamma < 1$ (Sublinear), discovery faces diminishing returns, perhaps due to human bottlenecks in reviewing AI-generated insights.
- If $\gamma = 1$ (Linear), discovery scales proportionally with compute, representing fully automated search.
- If $\gamma > 1$ (Superlinear), larger compute scales yield qualitative advantages[1] in intelligence, creating increasing returns to scale in discovery.

Conceptually, π describes the mapping from compute to effective search weight (accounting for congestion or quality tiers).

2.1 Population: The case for finite N

The model explicitly uses a finite population N. While much of the macroeconomic literature utilizes a continuum of agents for analytical tractability, a finite N is economically necessary to capture the "granular" nature of AI-driven growth. When AI agents scale economic production, a single person can run a billion dollar company.

The primary driver of wealth in this model is a series of discrete, winner-take-all breakthroughs. In a continuum limit, the probability of any single agent winning a discrete prize Π_t is zero, requiring one to model the prize as a measure-valued distribution and masking the essential economic reality of AI competition. Innovation is a high-stakes lottery where individual "lumpy" shocks do not average out, but rather drive aggregate inequality (see Gabaix (2011) on granular origins).

Second, the frontier AI sector is characterized by massive capital requirements and compute oligopolies. A continuum of atomistic, perfectly competitive searchers would be empirically inconsistent with the concentrated nature of agentic AI innovation. The finite N approach preserves the strategic density of this competition.

Furthermore, the transition thresholds for γ in the Propositions are not artifacts of the finite population. A continuum approach using the methodology of Luttmer (2007) would yield qualitatively identical regimes, as the fundamental tension between the search elasticity γ and the mean-reverting force of the wealth distribution remains invariant to the density of the measure. The finite N model simply allows us to track these dynamics without the unnecessary overhead of measure-valued Markov processes.

[1] The superlinear scaling case may also result if greater household wealth provides access to higher quality AI agents–that is, a non-competitive market for compute resources.

2.2 Prize scaling and expected returns

A central feature of the model is the assumption that the aggregate value of discovered capital, Π_t, equals the total search outlay: $\Pi_t = \sum_{j=1}^{N} m_{j,t}$. This specification can be interpreted as the competitive equilibrium of a frontier innovation sector with free entry. While stochastic at the individual level, the aggregate frontier obeys an expected value condition where breakthroughs require proportionally more compute resources to uncover. Even if a particular breakthrough is found with minimal individual search, the aggregate relationship reflects the macroeconomic reality that "nothing good is easy": the non-rival nature of AI-driven discoveries allows the winner to capture the capitalized value of rents proportional to the total market scale in expectation.

By equating the realized prize Π_t to the aggregate sum of m, the model assumes a competitive market structure in expectation where the total sunk costs required to push the technological frontier exactly exhaust the aggregate discovery rents. This scaling avoids the traditional "scale effects" critique (Romer, 1990; Aghion and Howitt, 1992) by ensuring that innovation incentives are not diluted by participation growth; rather, the rewards to innovation scale with the size of the reachable market.

Additionally, this assumption provides three critical analytical features:

- **Stationary Normalization:** Setting the expected aggregate return on innovation to zero ensures that in the forced-participation benchmark (Section 2.3), the aggregate capital stock K_t remains constant. This isolates the "pure" distributional effects of the technological parameter γ from confounding aggregate non-stationarity.

- **Robustness of the $\gamma = 1$ Threshold:** The core theoretical phase transitions at $\gamma = 1$ are robust to the scale of the prize. Under an alternative specification $\Pi_t = \zeta \sum m_{j,t}$ (where ζ represents sector productivity), the transition between ergodic ($\gamma < 1$) and condensation ($\gamma > 1$) regimes remains invariant. Because relative wealth shares are governed by the *ratio* of search efforts ($p_{i,t} \propto m_{i,t}^{\gamma}$), ζ merely scales the speed of accumulation without altering the elasticity threshold that triggers the winner-take-all feedback loop.

- **Consistent General Equilibrium Bridge:** In the endogenous search model (Section 2.4), the economy naturally gravitates toward an equilibrium where compute investment balances capital depreciation. This prize function ensures that "capital thinning"–the decline of the capital stock as $\gamma \to 0$–emerges strictly from the household's optimal response to declining search returns, rather than being hard-coded into the aggregate production function.

2.3 Forced lottery participation

To isolate the distributional role of stochastic innovation returns, the analysis first considers a benchmark economy in which households are required to allocate their entire capital income to the innovation lottery each period. This formulation preserves the aggregate capital stock while focusing strictly on the reallocation mechanics of the search process.

Specifically, given capital income $r_t k_{i,t}$, the forced participation benchmark imposes:

$$m_{i,t} = r_t k_{i,t}. \tag{4}$$

The lottery weight, or effective tickets, is determined by the mapping $\pi(m_{i,t}) = m_{i,t}^{\gamma} = (r_t k_{i,t})^{\gamma}$. Consequently, household i captures the innovation prize with probability:

$$p_{i,t} = \frac{(k_{i,t})^{\gamma}}{\sum_j (k_{j,t})^{\gamma}}. \tag{5}$$

Under this accounting, the aggregate prize equals the total capital income, $\Pi_t = r_t K_t$, and the household capital law of motion is:

$$k_{i,t+1} = (1-\delta)k_{i,t} + r_t K_t \mathbb{I}_{\{L_t = i\}}. \tag{6}$$

In this regime, capital accumulation is not driven by the endogenous choice of thrift, but by the stochastic "renewal" of depreciated capital through the innovation lottery.

Aggregate dynamics and steady state Summing the individual laws of motion (6) over all i yields the aggregate capital evolution. Since the lottery mechanism ensures exactly one winner per period ($\sum_i \mathbb{I}_{\{L_t=i\}} = 1$), the aggregate dynamics are deterministic:

$$K_{t+1} = (1-\delta)K_t + r_t K_t. \tag{7}$$

A stationary aggregate capital stock ($K_{t+1} = K_t$) implies that the rental rate must exactly offset the rate of obsolescence: $r^* = \delta$. Substituting the marginal product of capital $r_t = \alpha \bar{k}_t^{\alpha-1}$ (where $\bar{k} \equiv K/N$) into this condition pins down the unique steady-state capital per household:

$$\bar{k}^* = \left(\frac{\alpha}{\delta}\right)^{\frac{1}{1-\alpha}}. \tag{8}$$

This result provides a structural anchor for the model: the aggregate scale of the economy is determined by standard production parameters (α, δ), while the distributional morphology is governed entirely by the lottery elasticity γ. In steady state, the innovation prize

$\Pi^* = \delta K^*$ is equivalent to the total value lost through technological obsolescence. The lottery does not create aggregate capital; rather, it reallocates the renewal of the economy to a single winning innovator each period. Within this benchmark, the traditional role of saving is suppressed, allowing for a pure characterization of how wealth-contingent search access determines the long-run distribution.

Sublinear elasticity, $\gamma < 1$ When $\gamma < 1$, increases in household wealth translate into less-than-proportional increases in lottery access.

Proposition 1 (Ergodicity for $\gamma < 1$)**.** *Fix $N < \infty$. Suppose the lottery function satisfies*

$$\pi(m) = m^\gamma, \qquad 0 < \gamma < 1,$$

and let $\delta \in (0, 1)$. Then the Markov process governing the vector of household capital stocks

$$\mathbf{k}_t = (k_{1,t}, \ldots, k_{N,t}) \in \mathbb{R}_+^N$$

is positive Harris recurrent and aperiodic. Consequently, the process admits a unique invariant probability measure μ on $\mathbb{R}_+^N$, and for any initial condition $\mathbf{k}_0$ the distribution of $\mathbf{k}_t$ converges weakly to μ as $t \to \infty$.

Moreover, under μ the mean aggregate capital stock is finite and equals

$$\mathbb{E}_\mu \left[\sum_{i=1}^{N} k_i \right] = N\bar{k}^*,$$

where $\bar{k}^*$ is the unique solution to the aggregate steady-state condition

$$\delta \bar{k}^* = \alpha (\bar{k}^*)^{\alpha - 1}.$$

The formal proof is shown in Appendix B. In short, the proof uses $V(\mathbf{k}) = \sum_i k_i$ as a Lyapunov function and verifies a Foster-Lyapunov drift condition. Thus, when household capital is large, linear depreciation dominates the aggregate investment inflow. The condition $\gamma < 1$ ensures that no household captures a linear share of this inflow, preventing explosive concentration. Standard Markov chain stability results (Meyn and Tweedie, 1993) then imply the proposition.

The case $\gamma = 0$ corresponds to uniform lottery access and can be viewed as a limiting case of the sublinear regime. The technological limitation implies that the human household serves as a bottleneck to the innovation search. While this case eliminates any wealth-dependent reinforcement entirely, it shares the qualitative properties of the $0 < \gamma < 1$ case:

the induced wealth process is ergodic, admits a unique stationary distribution, and does not exhibit condensation.

Linear and superlinear elasticity, $\gamma \geq 1$ The steady-state distribution undergoes a regime change at the $\gamma = 1$ boundary, which separates mean-reverting and explosive wealth dynamics.

Remark 2 (Ergodicity)**.** *Throughout the paper, "ergodicity" is taken to mean positive Harris recurrence of the induced Markov chain for the vector of household capital stocks, implying existence of a unique invariant distribution and convergence in distribution from any initial condition. For* $0 < \gamma < 1$*, sublinear access to the lottery generates negative drift at high aggregate capital levels, ensuring ergodicity in this strong sense. By contrast, when* $\gamma = 1$*, proportional lottery access eliminates this drift: the individual wealth process becomes multiplicative, and positive Harris recurrence fails generically. In that case, stationary wealth distributions, when they exist, typically exhibit Pareto tails and need not possess finite moments.*

Proposition 2 (Critical proportional access, $\gamma = 1$)**.** *Suppose* $\gamma = 1$ *and* $\delta > 0$*. For any fixed* N*, the individual capital process of a household can be represented as a stochastic affine recursion*

$$k_{i,t+1} = M_t \, k_{i,t} + \varepsilon_{i,t},$$

where M_t *is a positive random multiplicative factor with* $\mathbb{E}[\log M_t] < 0$ *and* $\varepsilon_{i,t} \geq 0$ *is an idiosyncratic innovation capturing lottery wins.*

If the nondegeneracy condition

$$\mathbb{P}(M_t > 1) > 0$$

holds and the aggregate capital process admits a stationary distribution, then the stationary cross-sectional distribution of household wealth (if it exists) exhibits a Pareto upper tail. The tail exponent $\kappa > 0$ *satisfies the Kesten equation*

$$\mathbb{E}[M_t^{\kappa}] = 1.$$

Moreover, rank mobility is strictly lower than in the case $\gamma < 1$*, and wealth concentration is persistent.*

Superlinear elasticity, $\gamma > 1$ When the elasticity of access exceeds unity, the economy enters a regime of strong increasing returns to scale in discovery.

Proposition 3 (Condensation for $\gamma > 1$)**.** *Fix $N < \infty$ and $\delta \in (0,1)$. Suppose the lottery function satisfies $\pi(m) = m^\gamma$ with $\gamma > 1$. Then, for any initial condition $\mathbf{k}_0$, the capital distribution does not converge to a non-trivial stationary distribution. Instead, the economy exhibits* ***asymptotic condensation****:*

1. *The system leaves any compact set of non-degenerate distributions with probability one.*
2. *As $t \to \infty$, the capital share of a single household $s_{i,t} = k_{i,t}/K_t$ converges to 1 almost surely, while $s_{j,t} \to 0$ for all $j \neq i$.*
3. *The Markov chain is transient and admits no invariant probability measure on the interior of the simplex.*

The proof is provided in Appendix B. Intuitively, when $\gamma > 1$, the marginal return to search spending is increasing. Once a household acquires a sufficiently large share of wealth, their probability of winning the next innovation approaches certainty, creating an absorbing state of permanent dominance.

Tail regimes and condensation The qualitative regimes of the model are determined by the scaling of discovery access, γ. Figure 1 illustrates how the curvature of the lottery access function dictates the strength of the wealth-compute feedback loop.

Proposition 4 (Tail regimes)**.** *Fix $N < \infty$ and $\delta \in (0,1)$. The long-run morphology of the wealth distribution is governed by the elasticity of discovery access γ:*

1. ***Sublinear regime** ($0 \leq \gamma < 1$)**:** The capital process is positive Harris recurrent. The stationary distribution exhibits at most exponential tails, implying finite moments of all orders. Mean reversion dominates the stochastic reinforcement of wealth.*
2. ***Critical regime** ($\gamma = 1$)**:** Effective lottery weight is strictly proportional to spending. The individual capital process follows an affine-multiplicative recursion: $k_{i,t+1} = (1-\delta)k_{i,t} + \Pi_t \mathbb{I}_{\{L_t = i\}}$. Under the stationarity condition $\mathbb{E}[\ln(1-\delta+\Pi/k)] < 0$, the distribution exhibits a Pareto upper tail with exponent κ solving $\mathbb{E}[(1-\delta+\Pi/k)^\kappa] = 1$.*
3. ***Superlinear regime** ($\gamma > 1$)**:** A condensation phenomenon emerges. For any N, trajectories exist where a single household captures an $O(N)$ share of aggregate capital. This concentration occurs with probability one, the process is non-ergodic, and the system fails to admit a stationary distribution with finite mean.*

The economic intuition follows from the strength of the reinforcement loop. When $\gamma < 1$, wealth provides diminishing returns in the search for innovation, ensuring that shocks are

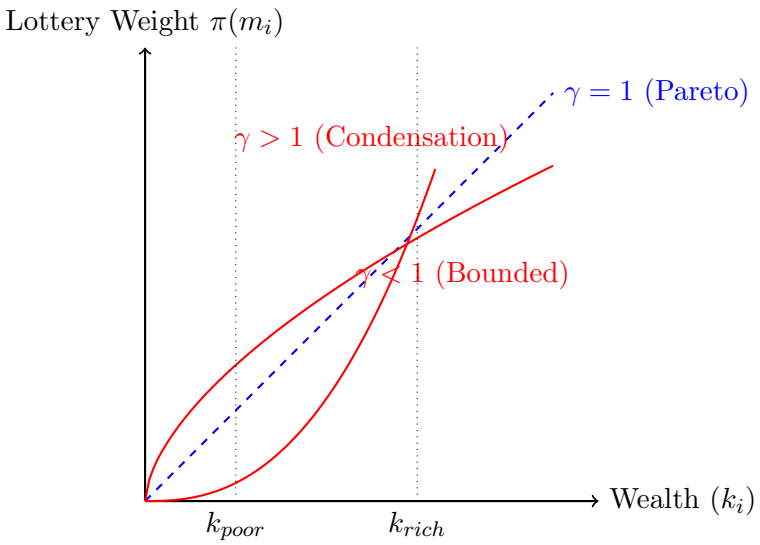


Figure 1: The Elasticity of Access to Innovation. The curvature of the discovery technology $\pi(m_i)$ governs the long-run morphology of wealth. Under sublinear access ($\gamma < 1$), diminishing returns to compute investment act as a structural bottleneck that restores mean reversion. Unitary scaling ($\gamma = 1$) recovers the classic random growth framework of Pareto-tailed distributions. Under superlinear access ($\gamma > 1$), the "Wealth buys Compute" feedback loop is amplified, triggering a phase transition toward oligarchic condensation.

temporary and the distribution remains compact. At the critical threshold $\gamma = 1$, the lottery exactly offsets linear depreciation, yielding the random growth foundation of classic Pareto laws. For $\gamma > 1$, the reinforcement is superlinear: the probability of winning the lottery grows faster than wealth, allowing a lucky innovator to "buy out" the competition. This leads to oligarchic condensation, where the long-run outcome is characterized by a persistent Dirac delta at the top of the distribution, implying vanishing long-run rank mobility.

2.4 Endogenous exposure: Households choose m

The preceding analysis assumed participation in the innovation lottery was a fixed fraction of wealth. This section endogenizes exposure by allowing households to choose their search intensity $m_{i,t}$. In standard precautionary savings models, risk aversion typically produces a "savings glut"; in the lottery economy, the variance of discovery creates a "search drought" where risk-averse agents may under-invest in the innovation frontier. Analyzing this environment establishes a benchmark for how search dynamics interact with traditional thrift in a general portfolio setting. The economy operates under a "technological dominance" regime where capital can be accumulated through two distinct channels:

1. **Innovation Search ($m_{i,t}$):** Investment in AI compute provides a stochastic probability $p_{i,t}$ of capturing the frontier prize Π_t. This capital represents the technological

state-of-the-art.

2. **Traditional Thrift ($s_{i,t}$):** Standard saving produces "legacy" capital. Legacy capital is assumed to be less compatible with the frontier: one unit of thrift generates only $(1-\delta_L)$ units of effective capital, where $\delta_L \geq \delta$ represents a higher rate of technological obsolescence.

The household maximizes $V(k_t) = \max_{m_t, s_t}\{u(c_t)+\beta\mathbb{E}[V(k_{t+1})]\}$ subject to $c_t+m_t+s_t = w_t + r_t k_t$ and the law of motion:

$$k_{t+1} = (1-\delta)k_t + (1-\delta_L)s_t + \mathbb{I}_{\{L_t=i\}}\Pi_t. \tag{9}$$

To ensure stationarity, households are assumed to have CRRA preferences $u(c) = \frac{c^{1-\phi}}{1-\phi}$. Defining detrended variables $\tilde{x}_t = x_t/G^t$, the problem becomes stationary in $\tilde{V}(\tilde{k}_t)$ with effective discount factor $\tilde{\beta} = \beta G^{1-\phi}$. Tildes are suppressed hereafter.

The optimal allocation between risk-free thrift (s_t) and risky search (m_t) is governed by:

$$\underbrace{(1-\delta_L)\mathbb{E}[V'(k_{t+1})]}_{\text{Marginal Value of Thrift}} \geq \underbrace{p'(m_t)\left[V(k_{t+1}^{win}) - V(k_{t+1}^{loss})\right]}_{\text{Marginal Value of Search}}. \tag{10}$$

When $\gamma < 1$, the marginal probability $p'(m_t)$ is strictly decreasing. Combined with DARA preferences, wealthy agents shift marginal savings into the risk-free asset s_t, reinforcing the mean-reversion of the ergodic regime. Conversely, for $\gamma > 1$, the marginal probability increases with spending. This creates an endogenous wealth threshold $\bar{\kappa}$ above which the wealthiest agents asymptotically abandon traditional thrift ($s_t \to 0$) to maximize lottery exposure.

Proceeding with the analytically cleaner baseline of immediate legacy obsolescence ($\delta_L = 1$), simplifies the portfolio to $s_t = 0$. In this case, the first-order condition for an interior optimum $m_i > 0$ is:

$$u'(c_t) = \beta\left\{\frac{\partial p_i}{\partial m_i}\mathbb{E}\big[V(k_{t+1}^{\text{win}}) - V(k_{t+1}^{\text{loss}})\big] + p_i\mathbb{E}\big[V'(k_{t+1}^{\text{win}})\big]\right\}, \tag{11}$$

where, defining the aggregate weight $T = \sum_j m_j^\gamma$, the marginal success probability is:

$$\frac{\partial p_i}{\partial m_i} = \gamma m_i^{\gamma-1}\frac{T - m_i^\gamma}{T^2}. \tag{12}$$

Remark 3 (The Risk-Aversion Stabilizer)**.** *With Decreasing Absolute Risk Aversion (DARA) preferences, the behavioral "cooling" of search intensity at the top of the wealth distribution*

acts as an endogenous stabilizer that reinforces the ergodic regime for low γ.

Standard heterogeneous agent models rely on precautionary savings to drive wealth accumulation. Here, the mechanism operates in reverse. Because the wage w_t provides a consumption floor, wealthier households become relatively more protective of their wealth as the lottery variance grows. This induces a concave search policy where $m^*(k)$ grows sublinearly, pulling the economy toward ergodicity unless technological convexity ($\gamma > 1$) explicitly overrides this behavioral smoothing.

Capital thinning. The introduction of the thrift channel $s_{i,t}$ clarifies the nature of the "search drought." Because risk-averse households prefer the certainty of legacy capital over the stochastic frontier, they socially under-invest in discovery. Even if $s_i > 0$ prevents a total collapse of the aggregate capital stock, a "capital thinning" effect persists through a shift in capital composition: the retreat into thrift starves the innovation frontier. As households shift from m_i to the less efficient legacy channel, the aggregate stock becomes dominated by older, higher-depreciation capital, resulting in a smaller productive base and higher rental rates r than a path of deterministic frontier accumulation.

Boundary Case of Zero Elasticity ($\gamma \to 0$) As $\gamma \to 0$, the discovery lottery success probability $P_i \to 1/N$ regardless of m_i. With strictly enforced $s_i = 0$, the marginal incentive to search on the intensive margin vanishes.[2] In the general case where $s_i > 0$ is available, $\gamma \to 0$ represents a "technological regression" boundary. The economy regresses from an AI-driven innovation regime to a traditional Solow-style thrift-based regime.

2.4.1 Linear Contest Case ($\gamma = 1$)

When $\gamma = 1$, the contest success function becomes linear:

$$p_i(m) = \frac{m_i}{\Pi}, \qquad \Pi \equiv \sum_{j=1}^{N} m_j. \tag{13}$$

A household with capital k chooses contest expenditure m subject to the budget constraint

$$c + m = w + rk.$$

[2]Specifically, $m = 0$ is not a Nash equilibrium. As $K \to 0$, the rental rate r and the shadow value $V'(k)$ approach infinity. A single household deviating from $m = 0$ to $m = \epsilon > 0$ captures the prize Π with certainty, and the infinite marginal product of capital at the zero-bound forces participation. universal participation then leads to a Bertrand-style race to the bottom on spending.

The continuation value satisfies

$$V(k) = \max_{m \in [0,\, w+rk]} \left\{ u(w + rk - m) + \beta \left[\frac{m}{\Pi} V(k + \Pi) + \left(1 - \frac{m}{\Pi}\right) V((1-\delta)k) \right] \right\}. \quad (14)$$

Importantly, households internalize the effect of their own effort on the prize:

$$\frac{d\Pi}{dm} = 1. \quad (15)$$

However, the marginal effect of m on winning probabilities remains linear, implying that the expected return to contest expenditure is independent of the scale of aggregate effort. As a result, the model does not pin down aggregate Π in decentralized equilibrium.

The first-order condition for an interior choice can be written as

$$u'(c) = \beta \left[\frac{1}{\Pi} \big(V(k + \Pi) - V((1-\delta)k) \big) + \frac{m}{\Pi} V'(k + \Pi) \right], \quad (16)$$

where the second term reflects the internalized effect of m on the prize itself.

Because aggregate contest effort is indeterminate, the decentralized equilibrium admits a continuum of allocations that differ only by a scaling of $\{m_i\}_{i=1}^N$. Imposing the feasibility normalization

$$\Pi = \delta K, \quad (17)$$

selects a steady-state which equates aggregate contest expenditure to aggregate depreciation and corresponds to the limit of the $\gamma < 1$ economy. Under this normalization, the equilibrium allocation coincides with the Ramsey steady state of a representative-agent economy. In particular, aggregate capital satisfies the standard Euler equation

$$1 = \beta \left(r + 1 - \delta \right), \quad (18)$$

and households consume part of their capital income each period. The $\gamma = 1$ case therefore represents a knife-edge economy in which individual effort choices are well defined, but aggregate contest activity must be fixed by an equilibrium selection rule.

The Ramsey-style normalization ($\Pi = \delta K$) prevents the aggregate economy from either evaporating or exploding infinitely. However, beneath that calm macroeconomic surface, individual wealth is growing and shrinking multiplicatively based on proportional lottery odds.

2.4.2 The Convex Regime and Oligarchic Condensation ($\gamma > 1$)

When the discovery process crosses into the convex regime ($\gamma > 1$), the search technology exhibits increasing marginal returns to scale. This phase transition fundamentally alters the nature of the equilibrium, rendering any symmetric or interior wealth distribution dynamically unstable.

Along the transition path, the economy is characterized by a preemptive Nash game. Because the marginal impact of compute expenditure, $\partial p_i / \partial m_i$, is increasing in m_i for low levels of investment, initial disparities in capital endowments are rapidly magnified. Agents engage in a war of attrition, over-investing to secure a probabilistic advantage. However, once a single agent secures a sufficient capital lead, the increasing returns heavily penalize the search efforts of the laggards, eventually driving the expenditures of the $N - 1$ trailing agents to zero.

Generically, the economy absorbs almost surely into a degenerate condensation limit: a single "oligarch" holds the entirety of the aggregate capital stock, while the remaining households consume only their wage endowment.

While this outcome represents the theoretical maximum of wealth inequality (with a Gini coefficient of $(N - 1)/N$), it restores aggregate capitalization to the efficient risk-pooled benchmark. In the condensation limit, the oligarch faces no viable competition, ensuring a winning probability of $p \to 1$. Consequently, the oligarch perfectly internalizes the aggregate law of motion, $K_{t+1} = (1 - \delta)K_t + m_t$. The household's recursive problem collapses identically to the standard deterministic Ramsey growth model.

The policy function in the convex regime exhibits a sharp discontinuity. For $k < K$, the marginal return to search is suppressed by the presence of the dominant agent, leading to a lockout equilibrium where $m(k) = 0$. Only at the boundary $k = K$ does the oligarch internalize the full return to capital, setting $m(K) = \delta K$. This discontinuity serves as the micro-foundation for the observed aggregate thinning: the shortfall in the capital stock is not due to a fundamental lack of productivity, but rather a strategic structure that prevents accumulation by any agent other than the incumbent.

3 Welfare

In the decentralized equilibrium, the lottery economy is subject to two fundamental inefficiencies. First, the winner-take-all nature of AI-mediated discovery imposes substantial idiosyncratic risk on households, leading to high wealth inequality or even oligarchic wealth condensation. Second, there is a search externality: an individual's compute expenditure increases the aggregate search intensity, reducing the winning probability for all

other agents. At the same time, higher compute spending accelerates the replenishment of the aggregate capital stock, as the innovation prize $\Pi_t = \sum_{i=1}^{N} m_{i,t}$ scales with the total resources allocated to the discovery frontier. The resulting welfare losses and inefficiencies generated by the decentralized search contest are quantified using a consumption-equivalent metric and a two-step decomposition that separates (i) the cost of idiosyncratic risk and (ii) the welfare consequences of distorted search incentives.

3.1 Decomposition of Welfare Gains

The quantification of welfare losses inherent in the decentralized search contest relies on an ex-ante perspective. In this framework, households evaluate expected lifetime utility behind a "veil of ignorance," prior to the realization of their initial wealth k_0 from the steady-state distribution μ^{DE}.

For a counterfactual policy regime CF characterized by the value function $V^{CF}(k_0)$, the ex-ante Consumption-Equivalent Variation (CEV), denoted by Θ, is defined as the uniform percentage scaling of the decentralized consumption path $\{c_t^{DE}\}_{t=0}^{\infty}$ required to equalize expected lifetime utility:

$$\mathbb{E}_{k_0 \sim \mu^{DE}} \left[\sum_{t=0}^{\infty} \beta^t u\big((1+\Theta) c_t^{DE}\big) \right] = \mathbb{E}_{k_0 \sim \mu^{DE}} \left[V^{CF}(k_0) \right]. \tag{19}$$

Under the assumption of CRRA utility with a coefficient of relative risk aversion $\phi \neq 1$, the scalar Θ admits the closed-form representation:

$$\Theta = \left(\frac{\mathbb{E}_{k_0 \sim \mu^{DE}}[V^{CF}(k_0)]}{\mathbb{E}_{k_0 \sim \mu^{DE}}[V^{DE}(k_0)]} \right)^{\frac{1}{1-\phi}} - 1. \tag{20}$$

This metric isolates the welfare consequences of the policy regime without requiring interpersonal utility comparisons. The total welfare gap is decomposed into two constituent counterfactuals to distinguish the utility cost of idiosyncratic risk from the efficiency costs of distorted search incentives.

(A) Pure Risk Removal with Fixed Allocations. To isolate the utility cost of lottery-based discovery, a counterfactual is constructed in which the stochastic contest is replaced by a deterministic return, while households are constrained to their decentralized compute outlays, $m^{DE}(k)$. Under this regime, each household i receives a deterministic capital increment equal to its compute expenditure. Since policy functions and aggregate resources remain fixed at decentralized levels[3], the resulting gap Θ^A captures strictly the insurance

[3] Consequently, prices r and w are fixed as well.

value of eliminating the variance inherent in the winner-take-all mechanism. Note that for $\gamma < 1$, the search success function is strictly concave, implying that the actuarially-fair expected return on search investment, $\mathbb{E}[\Pi]/m_i$, decreases with m_i. In an unequal equilibrium, this mechanism functions as an implicit redistribution from capital-rich to capital-poor agents. Consequently, Θ^A may take negative values if the welfare gains from risk-pooling are dominated by the loss of this implicit redistributive subsidy to borrowing-constrained agents at the bottom of the distribution.

(B) Full Efficiency Gain with Transitional Dynamics. The second counterfactual allows for the endogenous re-optimization of saving and compute behavior in a deterministic pooled environment. Because a risk-free steady state with identical preferences results in a degenerate, egalitarian wealth distribution, a pure steady-state comparison would conflate efficiency gains with mechanical redistribution. Instead, the analysis accounts for the full transitional dynamics. Upon an unexpected policy shift at $t = 0$, aggregates and prices evolve as the economy converges to the pooled steady state. Holding the initial distribution μ^{DE} fixed, Θ^B captures the net present value of the transition, accounting for both risk removal and the correction of aggregate investment incentives. This metric incorporates the "consumption sacrifice" required to shift the aggregate capital stock; a negative Θ^B indicates that the short-run costs of capital accumulation–discounted at rate β–outweigh the eventual long-run efficiency gains.

Computationally, these expectations are evaluated by taking the sample mean of value functions over the simulated cross-section $\{k_{i,T}\}_{i=1}^{N}$ in the decentralized steady state (for a large T).

Treatment of Non-Ergodic Regimes ($\gamma \geq 1$). In cases where the decentralized economy fails to admit a non-degenerate invariant distribution (e.g., condensation for $\gamma > 1$), the measure μ is replaced by the long-run empirical distribution derived from a large-scale finite-N simulation. Given a finite population, the economy achieves a quasi-stationary state where the winner's dominance is bounded by the population size, ensuring a well-defined ex-ante welfare calculation even in the presence of extreme wealth concentration.

3.2 Marginal Externality and the Pigouvian Wedge

The efficiency of the decentralized equilibrium can be characterized by comparing the household's marginal incentive to the marginal return to aggregate investment in a risk-pooled benchmark economy (Counterfactual B). In this counterfactual, the economy admits a deterministic aggregate law of motion for capital, as idiosyncratic discovery risk is fully

pooled. This approach allows for a rigorous efficiency analysis that remains invariant to the choice of social welfare weights.

3.2.1 Defining the Marginal Externality

Define the *Marginal Externality*, $\Delta(k)$, as the wedge between the social marginal benefit (SMB) and the private marginal benefit (PMB) of compute expenditure:

$$\Delta(k) = SMB(K) - PMB(k) \tag{21}$$

The Private Marginal Benefit, $PMB(k)$, is the return to search in the decentralized economy as derived from the household's recursive problem:

$$PMB(k) = \beta \left[\frac{\partial p(m;T)}{\partial m} \left(V(k^{win}) - V(k^{loss}) \right) + p(m;T) V'(k^{win}) \right] \tag{22}$$

where $T = \sum_j m_j^\gamma$ represents aggregate search intensity. The first term in the brackets captures the "prize-seeking" incentive–the marginal utility gain from increasing the probability of winning the lottery–while the second term captures the marginal utility of the prize expansion itself.

The Social Marginal Benefit, $SMB(K)$, is defined as the marginal continuation value of aggregate capital in the risk-pooled benchmark, rather than as a first-order condition of a planner's optimization problem.

The value function $V^*(K)$ captures the total return to investment in the risk-pooled benchmark economy, which by construction features no contest distortions or risk premia; it is a function of aggregate capital:

$$SMB(K) = \beta (V^*)'(K') \frac{\partial \Pi}{\partial m} = \beta (V^*)'(K') \tag{23}$$

where the prize-scaling identity $\partial \Pi / \partial m = 1$. The Marginal Externality $\Delta(k)$ thus measures the local gap in incentives (in utility units) caused by the stochastic nature of the discovery process.

3.2.2 The Implicit Pigouvian Wedge

This utility distortion maps into an interpretable policy metric by normalizing the gap by the marginal utility of consumption. This normalization yields the *Implicit Pigouvian Wedge*, $\nu(k)$, in units of the consumption good:

$$\nu(k) = \frac{\Delta(k)}{u'(c(k))} = \frac{SMB(K) - PMB(k)}{u'(c(k))} \tag{24}$$

This wedge represents the Pigouvian subsidy (if $\nu > 0$) or tax (if $\nu < 0$) required to align the household's marginal choice of compute with the efficient benchmark allocation. While the global welfare metrics (Θ^B) introduced in Section 3.1 integrate these costs over the transition path, the wedge $\nu(k)$ provides a local diagnostic of whether the economy is suffering from a "search drought" or an "arms race" at a specific wealth level.

3.2.3 Aggregate Wedge Statistics

The characterization of the systemic distortion inherent in the decentralized search contest is achieved by aggregating the state-dependent Pigouvian wedge $\nu(k)$ into two primary summary statistics. These measures capture distinct dimensions of the search externality across the wealth distribution.

The first statistic is the *Wealth-Weighted Aggregate Wedge*, ν_w, which quantifies the total capital-weighted distortion:

$$\nu_w = \int \nu(k) \frac{k}{K} d\mu(k). \tag{25}$$

This metric serves as a proxy for the aggregate efficiency gap. In regimes characterized by high wealth concentration, ν_w is dominated by the incentives of agents in the right tail of the distribution. A negative value for ν_w indicates that the majority of the economy's productive capital is controlled by agents engaged in a "preemption arms race," leading to aggregate over-investment in compute relative to the social optimum.

The second statistic is the *Median Search Wedge*, ν_m:

$$\nu_m = \text{Median}_\mu \left(\nu(k) \right). \tag{26}$$

While ν_w reflects distortions in aggregate capital renewal, ν_m characterizes the incentive structure facing the representative household. The divergence between ν_w and ν_m identifies structural heterogeneity in search incentives. For instance, in the sublinear regime ($\gamma < 1$), the coexistence of $\nu_m > 0$ (a risk-aversion-driven search drought for the median agent) and $\nu_w < 0$ (an arms race at the top) implies that no uniform policy can simultaneously align incentives for all agents.

By reporting both ν_w and ν_m, analysis distinguishes between the intensive margin of individual discovery and the aggregate efficiency of capital accumulation. While ν_w and ν_m provide a sufficient description of the efficiency-equity tradeoff, other moments of the distribution of $\nu(k)$ are informative in principle. A utilitarian-weighted mean, $\int \nu(k) d\mu(k)$, would assign equal weight to all households regardless of their capital contribution, reflecting a different normative weighting of the search externality. The choice of ν_w and ν_m in the

subsequent analysis prioritizes the distinction between aggregate production efficiency and the welfare of the median household.

3.2.4 Anatomy of the Distortion

The sign and magnitude of the implicit Pigouvian wedge $\nu(k)$ are determined by three mechanically distinct components:

$$\nu(k) = \nu_{\text{risk}}(k) + \nu_{\text{scale}}(k) + \nu_{\text{steal}}(k), \tag{27}$$

where the three terms are given by

$$\nu_{\text{risk}}(k) = \frac{\beta}{u'(c(k))} \left[(V^*)'(K') - V'(k^{\text{win}}) \right],$$

$$\nu_{\text{scale}}(k) = \frac{\beta}{u'(c(k))} \left[(1 - p(m;T))\, V'(k^{\text{win}}) \right],$$

$$\nu_{\text{steal}}(k) = -\frac{\beta}{u'(c(k))} \frac{\partial p(m;T)}{\partial m} \left[V(k^{\text{win}}) - V(k^{\text{loss}}) \right].$$

The risk term captures incomplete insurance across win/loss states; the scale term reflects failure to internalize aggregate prize expansion; the stealing term captures probability reallocation across agents. The risk and scaling terms are weakly positive whenever the social marginal return to capital exceeds the privately internalized marginal utility of winning, implying under-investment and thus a positive subsidy requirement. The stealing term is strictly negative because $\partial p/\partial m > 0$ and $V(k^{\text{win}}) > V(k^{\text{loss}})$, implying a private incentive to over-invest in order to appropriate winning probability from others.

The relative magnitude of these forces varies systematically with the discovery elasticity γ.[4]

Sublinear regime ($\gamma < 1$). Concavity of the search technology limits reinforcement and preserves competition across agents. The business-stealing term remains quantitatively important because marginal probability gains are not negligible relative to the scale term. As a result, ν_{steal} typically dominates in magnitude and $\nu(k)$ is negative over a large portion of the wealth distribution, implying private over-investment in compute relative to the social

[4]For $\pi(m) = m^\gamma$ and $T = \sum_j m_j^\gamma$, the marginal success probability is

$$\frac{\partial p_i}{\partial m_i} = \gamma m_i^{\gamma-1} \frac{T - m_i^\gamma}{T^2}.$$

Sublinearity ($\gamma < 1$) amplifies marginal probability gains at low spending levels through the factor $m^{\gamma-1}$, while superlinearity ($\gamma > 1$) concentrates marginal effectiveness among large spenders and drives $T - m_i^\gamma \to 0$ for the dominant agent in the condensation limit. These scalings govern the regime comparisons below.

benchmark.

Approaching the critical case ($\gamma \uparrow 1$). As γ increases toward unity, marginal effectiveness becomes more proportional to scale. High-wealth agents face stronger intensive-margin incentives, while low-wealth agents remain probability-constrained. The economy exhibits a wedge bifurcation: $\nu(k)$ becomes strongly negative in the upper tail (intensive-margin overinvestment), while the median distortion may remain small in magnitude. The aggregate distortion is therefore shaped by the divergence between ν_w and ν_m, reflecting asymmetric incentives across the distribution.

Critical case ($\gamma = 1$). With linear contest weights, $p_i = m_i / \sum_j m_j$ is homogeneous of degree zero in effort. Marginal probabilities and private marginal benefits are therefore well defined for interior effort profiles, and the wedge $\nu(k)$ is meaningful conditional on a selected decentralized steady state. However, aggregate effort is scale-indeterminate absent a normalization or equilibrium selection rule: only relative effort shares are pinned down. Under the selection of the Ramsey steady-state, $\nu(k) = 0$ for all k, implying that both wealth and compute distortions (ν_w, ν_m) vanish mechanically. This efficiency is a pointwise alignment of private and social incentives that exists only at this threshold. Once $\gamma > 1$, the linear proportionality fails and the emergence of increasing marginal returns in discovery triggers a preemption motive. Consequently, $\gamma = 1$ serves as a knife-edge result.

Convex regime ($\gamma > 1$). Superlinearity generates increasing marginal returns to effective search weight. Once a household acquires a sufficiently large capital share, its private marginal benefit converges to the social marginal benefit because it internalizes nearly the entire prize and the associated aggregate capital expansion. In the condensation limit, the dominant agent satisfies $\nu(k) \to 0$, while fringe agents have $\nu(k) > 0$ because both p and $\partial p / \partial m$ are negligible at low spending levels. Consequently, the wealth-weighted wedge ν_w collapses toward zero, while the median wedge ν_m becomes positive, reflecting under-investment by the majority of households excluded from the frontier.

Taken together, the three-term decomposition provides a unified account of the numerical patterns: sublinear economies are dominated by the stealing externality, the near-critical region features asymmetric intensive-margin distortions concentrated in the upper tail, and superlinear economies align incentives for the dominant agent while leaving the median household effectively disengaged from discovery.

4 Numerical results

5 Numerical results

Numerical simulation reveals sharp qualitative transitions in aggregate investment, wealth concentration, and mobility across the parameter space of γ. Rather than a smooth comparative static, the economy exhibits distinct regimes, each governed by a different structural mechanism of capital renewal and search intensity.

The analysis employs a benchmark parameterization: CRRA utility with risk aversion $\phi = 2.0$, discount factor $\beta = 0.95$, and a Cobb-Douglas production function with capital share $\alpha = 0.33$ and depreciation $\delta = 0.10$. The population size is fixed at $N = 200$. For each value of γ, the decentralized equilibrium is determined by iterating on optimal compute spending policies $m^*(k)$ and the consistent stationary aggregate capital stock K. In this framework, all stochasticity originates from the lottery mechanism; conditional on a realized discovery path, policy functions and welfare measures are deterministic. Detailed computational procedures are provided in Appendix C.

The welfare and distributional consequences of varying γ are summarized using the metrics established in Section 3. Wealth inequality is quantified by the Gini coefficient, while social mobility is measured by the Spearman rank correlation of wealth, ρ, over a 100-period horizon. Aggregate efficiency is assessed via *capital thinning*, defined as the percentage deviation of the decentralized capital stock from the forced-participation benchmark.

Welfare losses are reported in Consumption-Equivalent Variation (CEV) for the pure risk-pooling counterfactual (Θ^A) and the full-reoptimization transition (Θ^B). The structural distortion of incentives is characterized by the median Pigouvian wedge ν_m and the wealth-weighted wedge ν_w. Because ν_m captures the representative household's incentive while ν_w is dominated by the capital-weighted aggregate, the divergence between these measures identifies the tension between extensive-margin entry and intensive-margin accumulation at the top of the distribution.

Table 1 shows statistics for the baseline decentralized economy under various γ values. Table 2 shows statistics for the economy with varied ϕ and N values.

5.1 Regimes

The numerical results reveal four qualitatively distinct regimes. Although the numerical results impose $\delta_L = 1$ (so $s = 0$), the mechanisms should be interpreted as households reallocating toward legacy capital when frontier returns are low.

Table 1: Baseline Decentralized Economy Steady State Metrics

γ	Gini	ρ	K thinning (%)	CEV A (%)	CEV B (%)	ν_w (%)	ν_m (%)
0.05	0.943	0.258	99.8	9.9	227.5	-13.6	-71.9
0.10	0.921	0.351	99.5	9.6	178.9	-105.8	-100.4
0.15	0.919	0.437	99.1	9.9	140.5	-123.6	-102.8
0.20	0.915	0.426	98.5	9.1	112.3	-120.7	-93.0
0.25	0.914	0.411	97.5	7.3	91.1	-126.0	-99.6
0.30	0.910	0.543	96.3	6.5	75.5	-130.0	-101.2
0.35	0.916	0.503	94.3	4.6	59.0	-135.4	-99.9
0.40	0.934	0.533	92.1	3.9	48.7	-137.8	-96.0
0.45	0.932	0.610	88.9	2.9	38.2	-130.2	-98.7
0.50	0.940	0.485	84.6	2.0	29.7	-133.3	-98.2
0.55	0.953	0.631	78.9	1.2	22.0	-128.2	-95.0
0.60	0.948	0.648	73.8	0.8	16.9	-124.3	-95.5
0.65	0.961	0.753	66.5	0.5	11.5	-109.4	-97.8
0.70	0.965	0.688	54.7	0.1	4.3	-63.8	-90.1
0.75	0.981	0.882	48.4	-0.0	0.8	-2.4	-93.6
0.80	0.989	0.970	38.7	-0.1	-2.7	89.0	-97.2
0.85	0.995	1.000	27.8	-0.0	-3.3	645.8	-109.4
0.90	0.995	1.000	28.3	-0.0	-3.2	653.4	-105.1
0.95	0.995	1.000	56.9	-0.0	5.4	443.2	-231.7
1.00	0.995	1.000	46.8	-0.0	-0.3	0.0	0.0
1.05	0.995	0.000	46.8	0.0	-0.3	0.0	1.0
1.10	0.995	0.000	46.8	0.0	-0.3	0.0	1.0

Notes:
A. ρ reports the Spearman wealth rank correlation, which measures the extent to which a household's relative position in the wealth distribution remains constant over time by calculating the correlation between their wealth rankings in the numerical ergodic distribution (using 100 period separation).
B. $\nu > 0$ implies a Pigouvian subsidy (households under-invest) while $\nu < 0$ implies a Pigouvian tax (households over-invest).
C. For $\gamma = 1$, ν metrics are defined under the selection of the Ramsey steady-state.
D. For $\gamma > 1$, the capital-weighted wedge ν_w collapses to zero as the economy reaches the condensation limit; the sole dominant agent internalizes the aggregate prize expansion, effectively eliminating the search-induced stealing externality.

Regime I: Near-zero elasticity ($\gamma \approx 0$). When γ is close to zero, the optimal lottery investment $m(k)$ is nearly flat in wealth, so agents contribute similar amounts regardless of their capital holdings. This homogeneity does not imply low inequality. Instead, because aggregate investment is uniformly small, the lottery prize constitutes a large discrete shock relative to typical wealth levels. Moreover, winners do not rapidly dissipate these gains, since post-win incentives to scale investment are weak when $m(k)$ is insensitive to wealth. The resulting cross-section features a large mass of agents clustered at low wealth alongside a small number of persistent winners, which generates a high Gini coefficient even in the absence of aggressive or wealth-dependent bidding. The low marginal return to compute spending leads to severe *capital thinning* of frontier capital. CEV A is large and positive

Table 2: Robustness Variation: Metrics

$N = 200, \phi = 1.5$							
γ	Gini	ρ	K thinning (%)	CEV A (%)	CEV B (%)	ν_w (%)	ν_m (%)
0.05	0.928	0.283	99.9	13.7	293.5	-156.9	-104.5
0.10	0.912	0.363	99.6	16.1	206.5	-128.9	-100.3
0.25	0.916	0.424	98.2	14.7	109.3	-135.5	-100.8
0.50	0.937	0.469	90.5	6.4	43.4	-146.6	-97.9
0.75	0.979	0.856	58.9	0.3	6.6	-84.7	-92.3
1.00	0.995	1.000	46.8	-0.0	-0.3	0.0	0.0
1.05	0.995	0.000	46.8	0.0	-0.3	0.0	1.0
$N = 100, \phi = 2.0$							
γ	Gini	ρ	K thinning (%)	CEV A (%)	CEV B (%)	ν_w (%)	ν_m (%)
0.05	0.863	0.016	99.7	7.8	193.2	-113.5	-103.2
0.10	0.853	0.178	99.2	10.1	141.7	-111.5	-98.2
0.25	0.847	0.088	96.8	8.9	74.4	-119.9	-96.9
0.50	0.879	0.347	85.6	3.7	28.9	-124.5	-99.4
0.75	0.956	0.742	56.3	0.1	4.7	-60.4	-86.3
1.00	0.990	1.000	46.8	-0.0	-0.3	0.0	0.0
1.05	0.990	0.000	46.8	0.0	-0.2	0.0	1.9
$N = 500, \phi = 2.0$							
γ	Gini	ρ	K thinning (%)	CEV A (%)	CEV B (%)	ν_w (%)	ν_m (%)
0.05	0.975	0.680	99.9	5.0	282.6	452.0	-72.8
0.10	0.964	0.636	99.7	5.9	225.6	-75.6	-219.3
0.25	0.965	0.701	98.1	4.3	107.9	-157.3	-104.3
0.50	0.977	0.761	80.2	0.4	24.5	-161.4	-96.7
0.75	0.994	0.889	38.0	-0.0	-0.7	116.9	-96.9
1.00	0.998	1.000	46.8	-0.0	-0.4	0.0	0.0
1.05	0.998	0.000	46.8	0.0	-0.0	0.0	0.4

in this regime because lottery risk is severe relative to mean consumption, so eliminating idiosyncratic winner-take-all volatility provides substantial insurance value even when aggregate investment remains low. High CEV B values reflect the large long-run gains from re-optimizing investment and escaping a severely under-accumulated frontier capital stock, despite transition costs. The sign of the Pigouvian wedge on lottery investment depends on two competing scalings. Writing

$$PMB(k) \approx \frac{\beta}{T}\big(\gamma m^{\gamma-1}\Delta V + m^{\gamma}V'(k^{\text{win}})\big), \qquad SMB(K) = \beta(V^*)'(K'),$$

one sees that PMB carries a $1/T \sim 1/(N\bar{m}^{\gamma})$ factor while also containing $\gamma m^{\gamma-1}$, which becomes large as $m \downarrow 0$. Consequently, for very small γ and interior m the "steal" term $\gamma m^{\gamma-1}\Delta V/T$ can dominate and yield $PMB > SMB$ (so $\nu_m < 0$, private over-investment).

By contrast, increasing population size N raises T and suppresses PMB (scales like $1/N$) while $(V^*)'(K')$ remains of order unity; hence for large N, $\nu_m > 0$ (as seen in Table 2 for $N = 500$).

Regime II: Stable Churn. As γ increases, the marginal effectiveness of compute spending on AI search rises. Households begin to diversify their portfolios, allocating more resources to m. In this sublinear regime, diminishing returns to search ensure that wealth advantages are not self-reinforcing. High-wealth households face a lower marginal win probability per unit of wealth, leading to mean reversion. Rank persistence attains its minimum and there is a small dip in wealth inequality. Both Pigouvian wedges remain negative in this regime, indicating private over-investment in search relative to the social optimum, though distortions are modest compared to later regimes. CEV A declines as idiosyncratic volatility falls. CEV B remains positive, reflecting higher aggregate capital in the risk-pooled state even under discounted transition.

Regime III: The Competitive Arms-Race ($\gamma \to 1^-$). As the elasticity approaches unity from below, the marginal return to search spending increases sharply. Households are compelled to protect their market share against the business-stealing externality of competitors. Aggregate compute spending rises, and legacy investment s begins to be crowded out. A "wedge-splitting" occurs. The weighted wedge ν_w turns positive, while the median wedge ν_m remains negative, revealing the **intensive margin distortion**: the wealthiest agents, whose incentives dominate ν_w, over-invest in search to protect their rank against the business-stealing externality. Meanwhile, the median household continues to over-invest relative to the social benefit ($\nu_m < 0$), but its contribution is small relative to the intensive-margin escalation of top wealth holders. The arms-race is thus an elite phenomenon that increases aggregate capital but at the cost of high inequality. CEV A is close to zero in this regime as volatility is reduced by persistent dominance, while CEV B can turn negative: although capital (and consumption) in the risk-pooled steady state is higher, it is not high enough to offset the welfare costs of increased savings along the transition path to reach it.

Regime IV: Condensation ($\gamma > 1$). When γ exceeds one, the decentralized economy enters a condensation regime in which the lottery strongly amplifies existing wealth differences. Winning probabilities become sensitive to wealth. Once a household gains a sufficiently large wealth advantage, future wins become almost deterministic. As a result, the cross-sectional distribution collapses onto a single dominant household, while the remaining population holds negligible wealth, and total capital converges rapidly toward the

pooled steady state. For a non-leading household, PMB is effectively zero, but $SMB > 0$, hence $\nu_m > 0$. Intuitively, an incremental unit of non-leader effort has almost no effect on the household's winning chances but does raise the aggregate prize and thus social benefit. Meanwhile, $\nu_w = 0$, because the leader, as the de facto representative household, behaves optimally privately and socially. CEV A is zero as the single household faces no risk. CEV B is weakly negative in this regime because the transition to the pooled allocation requires imposing short-run consumption losses on the leader household despite higher long-run capital.

Robustness. The qualitative regime structure is robust to variation in population size N and preference curvature ϕ (Table 2). Rather than altering mechanisms, these parameters primarily affect the scale of inequality, welfare, and wedges in predictable ways:

- **Population size (N).** Because the innovation lottery is discrete, finite N mechanically bounds inequality: when exactly one household wins in each period, the Gini coefficient cannot exceed $(N-1)/N$. Reducing N therefore lowers the level of the Gini at low and moderate γ, reflecting the larger population mass represented by a single winner. Importantly, this scaling effect does not alter capital thinning, mobility patterns, or the location of the transition from mean-reverting to reinforcing wealth dynamics. Across N, the economy continues to exhibit stable churn for sublinear γ, an arms-race region as $\gamma \to 1^-$, and condensation at and above unit elasticity.

- **Risk aversion (ϕ).** Lower risk aversion increases households' willingness to finance compute, raising aggregate search effort and mitigating capital thinning at low γ. As γ approaches unity, lower ϕ amplifies competitive arms-race forces, increasing both aggregate investment and inequality. Consequently, welfare gains from pooling (CEV A and CEV B) and the magnitude of Pigouvian wedges vary with ϕ, but the qualitative distinction between the sublinear, arms-race, and condensation regimes remains unchanged.

- **Welfare and wedges.** The Pigouvian wedge ν_m varies systematically across regimes: it is typically negative at low γ (though not always), reflecting private over-investment driven by stealing incentives, becomes more negative in the arms-race region, and turns weakly positive only in the condensation regime where fringe agents under-invest. Meanwhile, ν_w becomes large and positive near $\gamma \uparrow 1$ and collapses in the condensation limit.

5.2 Policy interpretation

The regime classification has direct policy content because it maps technological primitives into distinct market distortions and thus into different policy instruments.

Low-γ (weak participation, misaligned incentives). When γ is low the principal inefficiency is insufficient participation in frontier search: private marginal returns are small relative to the social return and risk is high. Policy responses that expand baseline access to compute (subsidized access, broadened public provision of compute, or public–private pools that reduce idiosyncratic risk) raise aggregate investment and welfare while preserving incentives to innovate.

Churn (broadening access). In the stable-churn regime, redistributive or access-broadening policies that lower entry costs tend to increase contest participation and innovation. Policies that reduce fixed barriers or improve liquidity for smaller agents (e.g., access to shared cloud credits, standardized licensing of model checkpoints) can increase aggregate search effort and mobility.

Arms-race (coordination problem). Near the arms-race region, the economy faces a coordination problem: private incentives induce excessive defensive investment by the wealthy (intensive-margin waste), raising aggregate capital but lowering welfare net of distortion. The planner's problem is to discourage wasteful preemption while preserving productive search. Potential instruments include (i) non-price interventions that blunt the competitive externality (limits on exclusive model access, non-exclusive licensing requirements for publicly funded models), (ii) targeted subsidies or access for disadvantaged agents that increase effective competition, and (iii) regulation that raises the private cost of pure preemption without penalizing socially productive investment.

Condensation (structural reinforcement). For $\gamma \geq 1$, the inefficiency is structural: reinforcement dynamics concentrate rewards and collapse the extensive margin for the median agent. Traditional Pigouvian taxes are conceptually appropriate when the externality is local and observable, but they confront practical and political limitations in this domain (cross-border mobility of capital and compute, measurement and enforcement frictions, and potential regulatory arbitrage). Hence policy emphasis should shift toward instruments that change the structure of returns or access rather than trying to monitor and tax every instance of compute. Examples include (a) mandates on non-exclusive dissemination of frontier models from public funding, (b) institutional pooling of publicly subsidized compute

whose allocation is governed by access-enhancing rules, and (c) international coordination on standards for model sharing and licensing to avoid unilateral leakage.

Implementation trade-offs and political economy. All interventions carry trade-offs. Restrictions on exclusivity can reduce private appropriation but may blunt incentives for otherwise socially valuable private investment. Search pooling can raise output and mobility but concentrates governance power, whether in the hands of government officials (public provision of compute) or in a managerial class (private corporations). Moreover, many effective policies require international cooperation to prevent relocation of compute data-centers and other forms of arbitrage. Evaluating these trade-offs requires quantified, context-specific analysis of enforcement costs, leakage risks, and the political feasibility of coordination.

Understanding γ. Numerical results emphasize that policy choice depends on the underlying technology of access, γ. If high γ is the result of non-competitive access by the wealth elite to frontier AI models or more efficient compute, then pro-competitive policies are likely to be more effective. If, on the other hand, high γ represents increased capability of AI agents–that is, the reduction or elimination of the human-in-the-loop bottleneck–then probability of innovation lottery success arises directly from current household wealth and the ability to spend vast sums on AI agent compute. In this situation, strong re-distributive policies or policies of public AI compute spending may be required to prevent the economy from falling into an extreme oligarchic state.

6 Conclusion

The automation of innovation by AI agents alters the mechanics of capital accumulation, shifting the focus from incremental saving toward stochastic, search-driven discovery. In this environment, wealth inequality is not an artifact of heterogeneous preferences or entrepreneurial effort, but a structural consequence of the technological returns to scale in discovery, captured by the elasticity parameter γ.

The value of γ dictates a sharp transition across three distinct macroeconomic regimes that define the morphology of the wealth distribution:

- **The Ergodic Regime ($\gamma < 1$):** Diminishing returns to compute investment act as a structural bottleneck, restoring mean reversion and bounding long-run inequality.
- **The Pareto Regime ($\gamma = 1$):** Unitary scaling generates heavy-tailed distributions through multiplicative random growth, sustaining a "winner-take-most" equilibrium.

- **The Oligarchic Regime ($\gamma > 1$):** Superlinear returns trigger wealth condensation, causing the economy to collapse into a non-ergodic state of dynastic concentration.

Central to this framework is the *preemption motive*. Because existing wealth is the prerequisite for purchasing the compute necessary to enter the discovery lottery, a recursive feedback loop emerges: wealth purchases the probability of future wealth. This mechanism can override standard precautionary behavior, inducing even risk-averse agents to engage in discovery arms races that are socially excessive but privately rational for wealth preservation.

As AI assumes a greater role in economic production and the frontier of innovation, the determinants of social mobility increasingly depend on the *technology of search* rather than the *propensity to save*. This "Wealth of Winners" framework suggests that the technological returns to scale in discovery are a fundamental determinant of the long-run distribution. Consequently, effective policy design may require moving beyond traditional capital taxation to address the structural scaling of AI-mediated discovery and its implications for wealth condensation.

References

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.

Aghion, P., Akcigit, U., Bergeaud, A., Blundell, R., and Hemous, D. (2019). Innovation and top income inequality. *Review of Economic Studies*, 86(1):1–45.

Aghion, P. and Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, 60(2):323–351.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, 109(3):659–684.

Benhabib, J., Bisin, A., and Zhu, S. (2011). The distribution of wealth and capital income shocks. *Econometrica*, 79(1):123–174.

Bewley, T. (1977). The permanent income hypothesis: A theoretical formulation. *Journal of Economic Theory*, 16(2):252–292.

Cagetti, M. and De Nardi, M. (2006). Entrepreneurship, frictions, and wealth. *Journal of Political Economy*, 114(5):835–869.

Chodorow-Reich, G., Nenov, P., Santos, V., and Simsek, A. (2024). Stock market wealth and entrepreneurship. Technical Report 32643, NBER Working Paper.

Gabaix, X. (1999). Zipf's law for cities: An explanation. *Quarterly Journal of Economics*, 114(3):739–767.

Gabaix, X. (2011). The granular origins of business cycles. *Econometrica*, 79(3):733–772.

Goldie, C. M. (1991). Implicit renewal theory and tails of solutions of random equations. *Annals of Applied Probability*, 1:126–166.

Hubmer, J., Halvorsen, E., Salgado, S., and Ozkan, S. (2023). Why are the wealthiest so wealthy? new longitudinal evidence and implications for theories of wealth inequality. *Working Paper (SSRN)*. Conditionally accepted at *Econometrica*, available at SSRN: 4362818.

Jones, C. I. and Kim, J. (2018). A schumpeterian model of top income inequality. *Journal of Political Economy*, 126(5):1785–1826.

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

Korinek, A. (2024). Economic policy challenges for the age of ai. Technical Report 32980, NBER Working Paper.

Krapivsky, P. L. and Redner, S. (2001). Organization of growing random networks. *Physical Review E*, 63(6):066123.

Luttmer, E. G. J. (2007). Selection, growth, and the size distribution of firms. *American Economic Review*, 97(5):1647–1670.

Meyn, S. P. and Tweedie, R. L. (1993). *Markov Chains and Stochastic Stability*. Springer-Verlag, London.

Piketty, T. (2014). *Capital in the Twenty-First Century*. Harvard University Press.

Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy*, 98(5, Part 2):S71–S102.

Tullock, G. (1980). Efficient rent seeking. In Buchanan, J. M., Tollison, R. D., and Tullock, G., editors, *Toward a Theory of the Rent-Seeking Society*, pages 97–112. Texas A&M University Press, College Station.

Appendix

A Deterministic technological growth

Suppose aggregate output is given by

$$Y_t = A_t K_t^{\alpha} N^{1-\alpha}, \qquad A_{t+1} = (1+g)A_t,$$

with $g \geq 0$ deterministic. The rental rate of capital is

$$r_t = \alpha A_t (K_t/N)^{\alpha-1}.$$

Taking expectations over the lottery, aggregate capital evolves according to

$$\mathbb{E}[K_{t+1}] = (1-\delta)K_t + \Pi_t.$$

Under forced participation $\Pi_t = r_t K_t$, hence stationarity of the normalized aggregate requires $r_t = \delta$. Using $r = \alpha A_t (\bar{k}_t)^{\alpha-1}$ with $\bar{k}_t = K_t/N$, the balanced-growth condition reads

$$\delta = \alpha A_t (\bar{k}_t)^{\alpha-1},$$

so that

$$\bar{k}_t^* = \left(\frac{\alpha}{\delta}\right)^{1/(1-\alpha)} A_t^{1/(1-\alpha)}.$$

Define normalized capital

$$\tilde{k}_{i,t} \equiv \frac{k_{i,t}}{A_t^{1/(1-\alpha)}}, \qquad \tilde{K}_t = \sum_i \tilde{k}_{i,t}.$$

Substituting into the household law of motion shows that the normalized process $\{\tilde{k}_t\}$ obeys the same stochastic recursion as in the baseline economy with $A_t \equiv 1$. Consequently, all results for the trend-stationary distribution carry over verbatim to the normalized variables: for $0 < \gamma < 1$ the normalized process is positive Harris recurrent and admits a unique invariant distribution; for $\gamma = 1$ one finds Kesten-type Pareto tails; and for $\gamma > 1$ condensation occurs. Statements about inequality should therefore be interpreted as statements about capital holdings relative to the deterministic technology trend.

B Ergodicity and condensation under the model primitives

All results below assume the primitives stated in the main paper: a finite population of households $N < \infty$; Cobb-Douglas aggregate production with capital share $\alpha \in (0,1)$; depreciation $\delta \in (0,1)$; winner-take-all lottery that selects exactly one household per period; and the contest weight function $\pi(m) = m^\gamma$ with $\gamma \geq 0$. The period prize equals total search spending,

$$\Pi_t \;=\; \sum_{i=1}^{N} m_{i,t},$$

and the household law of motion is

$$k_{i,t+1} = (1-\delta)k_{i,t} + \Pi_t \cdot \mathbf{1}\{i \text{ wins at } t\}.$$

The finite-dimensional Markov chain $\mathbf{k}_t = (k_{1,t}, \ldots, k_{N,t}) \in \mathbb{R}_+^N$ is also used.

Remarks on boundedness of policies. In the forced-participation benchmark $m_{i,t} = r_t k_{i,t}$ and hence per-agent m is proportional to capital. In the endogenous-m economy, standard concavity and first-order condition arguments (and the budget constraint $c = w + rk - m$) imply that an optimal policy $m^*(k)$ is finite for finite k. It is natural and innocuous to restrict attention to policies taking values in a compact interval $[\underline{m}, \overline{m}]$ (this restriction can be relaxed with routine technical work; it is only used below to simplify constants in the drift/minorization arguments).

B.1 Lyapunov drift and existence of a stationary measure for $0 < \gamma < 1$

Consider a simple drift bound for aggregate capital that exploits the model's explicit expected evolution.

Lemma 1 (Aggregate drift)**.** *Define $V(K) = K$. Let $\overline{m}(K)$ denote the aggregate search spending implied by the decentralized policy at aggregate capital K. Under the model primitives there exist constants $M < \infty$ and $\lambda \in (0,1)$ such that for all $K \geq 0$,*

$$\mathbb{E}[K_{t+1} \mid K_t = K] - K \;\leq\; -\lambda K + B\mathbf{1}_{[0,M]}(K),$$

for some finite $B > 0$. In particular the drift is strictly negative for $K > M$.

Proof. From the household laws of motion and the definition of Π_t,

$$\mathbb{E}[K_{t+1} \mid K_t = K] = (1-\delta)K + \overline{m}(K),$$

so

$$\mathbb{E}[K_{t+1} \mid K_t = K] - K = -\delta K + \overline{m}(K).$$

Because policies are bounded on the compact grid used in computation (or because the FOC implies $m^*(k)$ grows at most linearly in k with a coefficient bounded by the rental rate), there exists $\overline{m} < \infty$ with $\overline{m}(K) \leq \overline{m}$ for all K. Choose

$$M = \frac{2\overline{m}}{\delta}, \qquad \lambda = \delta/2, \qquad B = \delta M + \overline{m},$$

so that for $K \geq M$, $\overline{m}(K) \leq \frac{\delta}{2}K$ and hence $\mathbb{E}[K_{t+1} \mid K_t = K] - K \leq -\lambda K$. For $K \leq M$ the left-hand side is bounded above by the finite constant B. This yields the stated drift inequality. □

Combining the drift bound with a small-set (minorization) condition gives recurrence. The minorization below uses only the model's discrete lottery randomness plus continuity of the deterministic update maps.

B.2 Minorization (small set) via finite winner sequences

Lemma 2 (Minorization on a compact set)**.** *Fix the compact set $C = \{K : 0 \leq K \leq M\}$ with M from Lemma 1. There exist an integer $n \geq 1$, $\eta > 0$, and a probability measure ν on $\mathbb{R}_+^N$ such that for every initial $\mathbf{k}$ with aggregate $K \in C$ and measurable $B \subset \mathbb{R}_+^N$,*

$$\mathbb{P}\big(\mathbf{k}_{t+n} \in B \mid \mathbf{k}_t = \mathbf{k}\big) \;\geq\; \eta\, \nu(B).$$

Thus C is a small (petite) set.

Proof. The key observation is that the only source of randomness is the winner identity at each date; this is a finite-space random draw that assigns positive probability to any given winner sequence of length n. Fix a target compact neighborhood $U \subset \mathbb{R}_+^N$ of some desirable interior point $\mathbf{k}^0$ (for example a small ball around the deterministic vector $(m_i/\delta)_{i=1}^N$). Consider the deterministic update map

$$\Phi_{(i_1,\ldots,i_n)}(\mathbf{k}) \;=\; \mathbf{k}_{t+n}$$

obtained by applying the winner sequence $i_1, \ldots, i_n$ deterministically (i.e., treat the randomness as a control sequence). This map is continuous in $\mathbf{k}$ because the per-period update is continuous in the inputs (given the prize and the deterministic policy mapping). By continuity, there exists a neighborhood $U_{\mathbf{k}}$ of $\mathbf{k}$ such that $\Phi_{(i_1,\ldots,i_n)}(U_{\mathbf{k}}) \subset U$.

Because each winner sequence $(i_1, \ldots, i_n)$ occurs with strictly positive probability (product of per-period winner probabilities conditional on the path), the n-step kernel from any $\mathbf{k} \in C$ assigns at least some positive mass to U. Cover the compact set C by finitely many such neighborhoods $\{U_{\mathbf{k}_j}\}_{j=1}^{J}$. Let $\eta > 0$ be the minimum over the finite cover of the positive probability that the corresponding winner sequence drives the state into the chosen target neighborhood U. Define ν to be the probability measure uniform over U (normalized Lebesgue on U). Then for all $\mathbf{k} \in C$ and measurable B,

$$\mathbb{P}(\mathbf{k}_{t+n} \in B \mid \mathbf{k}_t = \mathbf{k}) \geq \eta\, \nu(B).$$

This is precisely the minorization condition required. (The argument is constructive: one can choose n large enough so that some finite sequence of wins leads into U from every neighborhood in the finite cover.) □

B.3 Positive Harris recurrence for $0 < \gamma < 1$

Proposition 5. *Suppose $0 < \gamma < 1$ and the model primitives above hold. Then the Markov chain $\{\mathbf{k}_t\}$ on $\mathbb{R}_+^N$ is positive Harris recurrent and admits a unique invariant probability measure. Moreover, the chain is aperiodic and ergodic.*

Proof. Lemma 1 provides a Foster-Lyapunov drift inequality with small set C. Lemma 2 supplies the minorization (petite) condition on C. The combination of a drift condition and a petite set implies positive Harris recurrence and geometric ergodicity on compacts by standard results (see Theorem 15.0.1 and related results in Meyn and Tweedie (1993)). Aperiodicity is immediate because there is always positive probability that the same household wins in two consecutive periods, preventing periodic cycles of deterministic length. This completes the proof. □

B.4 Condensation in the superlinear regime $\gamma > 1$

Formalizing the condensation claim used in the text uses the share dynamics derived from the household law of motion.

Proposition 6 (Condensation for $\gamma > 1$)**.** *Fix $N < \infty$, $\delta \in (0, 1)$, and suppose $\pi(m) = m^\gamma$ with $\gamma > 1$. For any interior initial condition (all $k_{i,0} > 0$), with probability one there exists a (finite but random) time T and an index i^* such that*

$$\lim_{t \to \infty} s_{i^*,t} = 1,$$

where $s_{i,t} = k_{i,t}/K_t$. In words: almost-sure condensation (one agent eventually holds essentially all aggregate capital) occurs.

Proof. The proof proceeds in three steps.

Step 1: Local amplification. Write shares $s_{i,t} = k_{i,t}/K_t$. Conditional on agent i winning at t, the share update is

$$s_{i,t+1} = \frac{(1-\delta)s_{i,t} + \Pi_t/K_t}{(1-\delta) + \Pi_t/K_t} = \frac{(1-\delta)s_{i,t} + \overline{m}_t/K_t}{(1-\delta) + \overline{m}_t/K_t},$$

where $\overline{m}_t = \Pi_t$. Because $\pi(m) = m^\gamma$ with $\gamma > 1$, when a share $s_{i,t}$ is large the win probability $p_{i,t}$ is asymptotically close to one and the marginal amplification of share upon winning is superlinear. Concretely, there exists a threshold $s^\star \in (0,1)$ and a constant $\eta > 0$ such that whenever $s_{i,t} \geq s^\star$ the conditional expected increment of $s_{i,t}$ given a win is at least $\eta > 0$, and the conditional probability of winning satisfies $p_{i,t} \geq 1-\varepsilon$ with ε arbitrarily small for sufficiently large $s_{i,t}$.

Step 2: Formal bounding of the transition probabilities. Fix $\varepsilon \in (0, 1/2)$ and let $s_{\text{high}} = 1 - \varepsilon$. Suppose agent i achieves a share $s_{i,t} \geq 1 - \varepsilon$. The remaining agents share at most ε. Because the contest success function $\pi(m) = m^\gamma$ is strictly convex for $\gamma > 1$, the sum of the competing weights is maximized when a single competitor holds the entire remaining share. Thus, $\sum_{j \neq i} m_{j,t}^\gamma \leq \varepsilon^\gamma$. Agent i's probability of winning the next period is therefore bounded below by:

$$p_{i,t} \geq \frac{(1-\varepsilon)^\gamma}{(1-\varepsilon)^\gamma + \varepsilon^\gamma} \tag{28}$$

The probability that the leader loses is $1 - p_{i,t} \leq \frac{\varepsilon^\gamma}{(1-\varepsilon)^\gamma + \varepsilon^\gamma}$. Since $\varepsilon < 1/2$, $(1-\varepsilon)^\gamma > (1/2)^\gamma$. The probability of a loss is bounded by $1 - p_{i,t} \leq C\varepsilon^\gamma$, where $C = 2^\gamma$. Because $\gamma > 1$, the probability of losing the lottery shrinks superlinearly relative to the distance to the absorbing state ε.

Step 3: Almost-sure hitting time via Borel-Cantelli. To show that the high band $[s_{\text{high}}, 1)$ is reached almost surely, let $S_t = \max_i s_{i,t}$. Since the state space is a finite simplex and the Markov chain is irreducible on the interior prior to condensation, there exists a uniform lower bound $\eta > 0$ and an integer $L \geq 1$ such that from any initial distribution of wealth, the probability of a single agent winning L consecutive times is at least η. Partition time into discrete blocks of length L. Let E_m be the event that a single agent wins all lotteries in the m-th block. Because $\mathbb{P}(E_m) \geq \eta > 0$ independent of prior history,

$\sum_{m=1}^{\infty} \mathbb{P}(E_m) = \infty$. By the second Borel-Cantelli lemma, the event E_m occurs infinitely often almost surely. For a sufficiently large but finite L, consecutive wins guarantee the winner's share exceeds s_{high}. Once in this absorbing band, Step 2 ensures the supermartingale properties drive the share to 1 almost surely, completing the proof of condensation. □

Remarks. The condensation argument above is tailored to the finite-N winner-take-all environment studied in the paper. It relies on the discrete randomization of winner identity (so that favorable win sequences have strictly positive probability) and on the superlinear reinforcement $\gamma > 1$. The proof gives a constructive route to the almost-sure condensation statement used in the main text; stronger limit theorems (for example on the rate at which dominance occurs) may be derived under additional assumptions on initial conditions and on the behavior of $m^*(k)$ for large k.

B.5 The critical case $\gamma = 1$ and Kesten tails

When $\gamma = 1$ the individual recursion can be written in affine-multiplicative form $k_{i,t+1} = M_t k_{i,t} + Q_{i,t}$ where the multiplicative factor M_t is random and captures the aggregate-scale reinforcement of winners. Standard Kesten theory applies: under the mild nondegeneracy conditions stated in the main text (e.g., $\mathbb{P}(M_t > 1) > 0$ and $\mathbb{E}[\log M_t] < 0$ when conditioning on aggregate stationarity), any stationary cross-sectional distribution has a Pareto upper tail with exponent $\kappa > 0$ determined by $\mathbb{E}[M_t^{\kappa}] = 1$. Refer to Kesten (1973) and Goldie (1991) for the formal conditions and proofs.

C Computational Implementation

Key implementation choices and algorithms are summarized below; the paper's main text describes the economic primitives and objects being computed.

Grids, choice set and VFI. Discretize individual capital on a log-spaced grid with support wide enough to accommodate both egalitarian and highly concentrated outcomes. Candidate compute choices are taken on a grid (including zero) of log-spaced fractions of current income. For each conjectured aggregate capital stock K (and implied prices (r, w) and prize $\Pi = \delta K$), the household problem is solved by value-function iteration (VFI) on the discrete k-grid, using linear interpolation to evaluate continuation values off-grid.

Equilibrium fixed point and simulation. For a fixed γ, the decentralized stationary equilibrium is computed via a nested fixed-point algorithm. The inner loop solves the household problem for a given conjecture of the aggregate contest mass T, yielding an

optimal policy $m^*(k; K, T)$. The outer loop updates the aggregate capital stock K until the market-clearing condition $\int m^*(k)\, d\mu^{DE}(k) = \delta K$ is satisfied. The stationary distribution μ^{DE} is approximated by simulating a cross-section of agents (with fixed random seeds for reproducibility), using a long burn-in followed by a long sample window to compute ergodic averages. Beliefs about the aggregate contest mass are updated iteratively until convergence, and the outer fixed point is obtained using a bracketing root-finding method.

CEV computation. All ex-ante lifetime utilities are evaluated by averaging individual value functions over the simulated decentralized cross-section $\{k_i\}_{i=1}^N$. Counterfactual A replaces the stochastic prize with its deterministic counterpart while holding the decentralized policy $m^{DE}(k)$ fixed; the resulting deterministic problem is solved by VFI on the same k-grid. Counterfactual B is computed along the deterministic transition path: starting from the decentralized cross-section, iterate forward for T_{trans} periods, recomputing prices (r, w), pooled optimal policies, and value functions at each date. Individual capital evolves according to $k_{i,t+1} = (1-\delta)k_{i,t} + m_t(k_{i,t})$; discounted flow utility is accumulated along the transition, and the final period's pooled value function provides the terminal continuation value. Consumption-equivalent variations (CEVs) are computed using the closed-form CRRA transformation.

Pigouvian wedge $\nu(k)$. The private marginal benefit (PMB) is computed pointwise on the k-grid from the household first-order condition:

$$PMB(k) = \beta\Big(\tfrac{\partial p}{\partial m}\,[V(k^{\text{win}}) - V(k^{\text{loss}})] \;+\; p(m;T)\,V'(k^{\text{win}})\Big),$$

where $V'(\cdot)$ is obtained via centered finite differences of the decentralized value function and continuation values are evaluated by linear interpolation. The social marginal benefit (SMB) is computed as a scalar by numerically differentiating average lifetime utility in the pooled deterministic benchmark with respect to aggregate capital: solve the pooled problem at K and $K+\varepsilon$, holding the initial cross-section fixed, and form

$$SMB \approx \beta \frac{\overline{V^*}(K+\varepsilon) - \overline{V^*}(K)}{\varepsilon}.$$

The implicit Pigouvian wedge is then $\nu(k) = (SMB - PMB(k))/u'(c(k))$. Statistics are calculated by standard aggregation (e.g. mean, median, etc.).

Numerical safeguards and non-ergodic cases. To avoid spurious extrapolation, next-period states are clipped to the k-grid and contest denominators are bounded away from zero. When the decentralized process does not admit an analytic invariant distribution (e.g.

under condensation for $\gamma \geq 1$), ex-ante welfare and wedge statistics are computed using the long-run empirical cross-section generated by the simulation.