Financial Information Technology and the Inequality Gap

Roxana Mihet[‡]
Swiss Finance Institute at HEC Lausanne

August 27, 2025

Abstract

Who benefits from financial information technology depends on which friction falls. In a baseline noisy rational-expectations model with fixed participation and information costs, lowering participation costs broadens entry, improves risk sharing, and reduces wealth and return inequality. Lowering information costs reallocates surplus to informed traders, concentrates intermediation, and increases cross-sectional inequality. The model delivers clean distributional predictions: participation, volumes, premia, and fee incidence, that align with recent trends. To share gains widely, technology needs to target participation frictions. Subsidizing data and promoting pooling funds exacerbates inequality.

Keywords: Financial information technology, equity markets, inequality, Gini, delegated wealth management.
JEL Classification: E21, G11, G14, L1, L15

[*]Contact: University of Lausanne, Quartier Chamberonne, Extranef 238, Lausanne, Vaud 1016, Switzerland. Email: roxana.mihet@unil.ch.

[†]First draft: November 2019. This draft: August 2025. We are grateful to Anna Bayona (discussant), Jess Benhabib, Stijn Claessens, Jerome Dugast (discussant), Andreas Fuster, Mark Gertler, Avi Goldfarb, Paymon Khorrami, Luc Laeven, Andrew Lo, Joseba Martinez, Guido Menzio, Emily Mochini, John Muellbauer, Cecilia Parlatore, Luigi Pistaferri, Thomas Philippon, Claudia Sahm, Thomas Sargent, Lin Shen (discussant), Hyun Shin, Johannes Stroebel, Laura Veldkamp, Venky Venkateswaran, Xavier Vives, and Ansgar Walther for their useful suggestions, as well as participants at numerous conferences and workshops for their helpful comments. This paper won the Best Dissertation Award in the Inequality Category at NYU, and the ECB Young Economist Best Paper Award. This research was funded by the BFI Macro-Financial Modeling Dissertation Grant, NYU Stern, and The Sandoz Foundation - Monique de Meuron Program for Academic Excellence. All errors are our own.

1 Introduction

The world is facing an increasingly problematic gap in capital income and wealth inequality. FinTech and Artificial Intelligence have the potential to narrow this gap by lowering the costs of market access and by democratizing tools for information processing (e.g., Duffie et al. (2022), Foucault et al. (2025)). We see platforms expanding access to basic financial services, credit, and diversified investments for households that previously lacked them (e.g., D'Acunto and Rossi (2021), D'Acunto et al. (2019), Fuster et al. (2022), Reher and Sokolinski (2024)). At the same time, market-intelligence providers now aggregate and clean an ever-growing universe of data—traditional feeds: prices, analysts' reports, news , and alternative sources: social media, satellite imagery, geolocation, transaction traces (e.g. Grennan and Michaely (2021); Katona et al. (2025)). But a glut of data does not automatically translate into useful information (e.g., Dugast and Foucault (2025); Dessaint et al. (2024)). Extracting value requires algorithms, organization, and skill. Innovations in financial information technology therefore push on two levers at once: they make participation simpler and cheaper, and they change the economics of becoming and being an informed investor.

We study how those two levers shape the capital wealth distribution. We develop a baseline noisy rational-expectations model in the tradition of Grossman and Stiglitz (1980) and Verrecchia (1982) with investors trading in a risky equity asset and a safe bond, with a portfolio-choice participation margin emphasized by Peress (2005). An investor chooses whether to stay out of the equity market, to participate uninformed, or to participate and self-acquire private information. Technology shifts enter through three conceptually distinct costs: a fixed participation friction (i.e., onboarding, basic literacy, taxes/admin knowledge), an adoption cost of becoming informed (i.e., the set-up required to turn data into useful signals), and an information intensity cost that governs how precise signals are once informed. Equilibrium is pinned down by a fixed point in the share of informed and participating investors.[1] Prices aggregate fundamentals and noise as in the classic framework. Equilibrium price informativeness rises when more investors are informed or when informed investors raise precision. We track distributional implications through the Gini, Lorenz curves, and price informativeness as the three objects of interest.

We evaluate technology counterfactuals under two canonical wealth distributions.

[1] Existence (and, under mild curvature, uniqueness) follows from standard monotonicity arguments, and we compute it with a simple contraction mapping. Formal statements and the algorithm are in Sections 2.3–2.4 and the Appendix.

A Lognormal captures the bulk of the distribution where most households live; a Pareto captures the empirically heavy upper tail that macro-inequality work emphasizes (e.g., Benhabib and Bisin (2016); Gomez (2023)). Using both lets us separate forces that act "in the middle" from those that load on the tail. It also guards against functional-form critiques: our mechanisms should not depend on a single parametric shape. Intuitively, the Lognormal concentrates mass near the median, so shifts that bring middle-wealth households into the market show up strongly, whereas the Pareto puts substantial mass in the right tail, exactly where the adoption and intensity margins bite hardest. Our emphasis on distributional statistics aligns with recent work on household finance and inequality (e.g., Badarinza et al. (2016); Campbell et al. (2019)).

The model delivers two simple components that together move the Gini and the Lorenz curve: (i) *Between-group differences*: informed vs. uninformed participants earn different risk-adjusted returns because informed investors time and size risk better; and (ii) *Within-group dispersion*: among the informed, precision scales with effective resources/benefits, so information intensity stretches outcomes even holding group shares fixed.

Technology that cuts the participation cost mainly expands the set of uninformed participants drawn from the middle of the distribution. Under a Lognormal, this is a big mass: the Lorenz curve shifts up and the Gini falls, because many households newly access the equity premium and the between-group gap shrinks (a larger share of the market earns roughly uninformed returns). Under a Pareto, the same effect is present but smaller—the tail still dominates overall inequality.

Cutting the information adoption cost moves investors from the uninformed group into the informed group. Under a Lognormal, that reallocation comes mostly from the middle, raising the between-group component as more households cross the information threshold and earn higher risk-adjusted returns. The Lorenz curve bends down in the upper quantiles and the Gini rises. Under a Pareto distribution, the effect is stronger: the right tail contains many investors close to or above the adoption threshold, so the informed share jumps disproportionately at the top, thickening the tail and raising inequality even more.

Cutting the information intensity cost largely leaves group shares unchanged, but increases precision among those already informed. That works mainly through within-informed dispersion. In a Lognormal world, this has a modest effect on the Gini: the group is relatively small and the middle remains uninformed. In a Pareto world, it matters more: the informed are concentrated at the top, so raising their precision shifts mass further to the right tail. In both cases, price informativeness rises; whether the

Gini ticks up or stays flat depends on how strongly the precision gains feed through to realized cross-sectional returns relative to the general-equilibrium dampening of gains as prices become more informative.

A natural critique is that these patterns might be "entry-driven." They are not. When the participation friction is set to zero so that everyone participates, cheaper adoption still reallocates mass from uninformed to informed (raising the between-group component), and cheaper intensity still stretches outcomes within the informed group. The Gini and Lorenz implications above are preserved, with larger movements under Pareto because the tail is thicker.

Crucially, the shape of information costs matters. When the information adoption cost includes a fixed component, the model features a clear threshold: agents either remain uninformed or pay the set-up cost and become informed. In that case, technology that lowers adoption has a strong between-group effect. By contrast, when the fixed adoption component is absent (purely marginal costs of precision), the extensive "become informed" margin collapses into a continuous precision choice. Some investors then choose minimal precision rather than staying strictly uninformed. Inequality movements are driven much more by within-informed dispersion (via intensity costs) and much less by between-group reallocation, with the effect again amplified under a Pareto tail. We highlight both cases in the analysis and show how the Lorenz curves and price informativeness respond under each cost shape.

Because most investors access research through intermediaries, we add a concise extension with delegation. Investors can search for a fund and share information acquisition costs through a fund manager (see Gârleanu and Pedersen (2018); Sockin and Xiaolan (2023); Greig et al. (2024); Bonelli and Foucault (2024) for complementary evidence on managerial advice.). Pooling creates scale economies: some households that would not self-learn optimally can become informed by delegating, while the wealthiest still prefer to self-acquire precision and the least wealthy remain uninformed. We then study technology that lowers delegation frictions: easier search and screening of managers, better data infrastructure in asset management, and lower marginal costs of research at scale. Lower search frictions expand the delegated sector by drawing mass from uninformed into delegated informed, compressing the between-group component and improving inclusion; lower marginal research costs raise the precision of delegated information, which can increase within-informed dispersion if self-learners do not keep pace. Our results connect to active/passive interactions and manager selection (e.g. Gârleanu and Pedersen (2022) , Berk and Green (2004), Berk and van Binsbergen (2016)). Overall inequality reflects the balance of these forces, with

stronger compression when the middle of the distribution moves into delegated informed (Lognormal) and stronger stretching when improvements load on the top tail (Pareto). Throughout, price informativeness rises as the stock of informed precision increases; general-equilibrium feedback tempers private gains as prices reveal more.

Our main contribution is to connect classic noisy rational-expectations models (Grossman and Stiglitz, 1980; Verrecchia, 1982; Admati and Pfleiderer, 1990; Peress, 2004) to modern FinTech/AI and macro-inequality themes (e.g., Acemoglu and Restrepo (2017), Korinek and Stiglitz (2018) Acemoglu (2024), Acemoglu (2025), Autor (2024) on the distributional consequences of AI, Benhabib et al. (2017), Benhabib et al. (2019), Gomez (2023), Gaillard et al. (2023) on tails and heterogeneous returns; Grennan and Michaely (2021), Rossi and Utkus (2024), Greig et al. (2024), Katona et al. (2025) on FinTech and household finance). Our contribution is to (i) separate participation, adoption, and intensity costs;[2] (ii) map each to between-group and within-informed components of inequality; (iii) perform comparative statics exercises under Lognormal and Pareto initial wealth distributions and numerically document movements in the Gini, Lorenz curves, and final wealth thresholds and distributions; and (iv) include a minimal delegation extension that lines up with how information is purchased in practice, without being essential for the baseline results.

The paper proceeds as follows. Section 2 lays down the baseline economy, characterizes the competitive rational-expectations equilibrium, and studies comparative statics theoretical exercises with respect to technology that lowers (i) participation frictions and (ii) the costs of adopting and intensifying information. Section 3 numerically quantifies the implications for inequality: Gini indices, Lorenz curves, and final wealth distributions under Lognormal and Pareto initial wealth distributions. Section 4 presents a variety of robustness exercises, discusses the relation to the empirical evidence, and presents theoretical generalizations of the baseline model. Section 5 extends the framework to delegation via funds and analyzes how falling intermediation costs affect inclusion and inequality. Crucially, the results of the numerical quantification mirror the baseline economy without delegation. Section 6 discusses policy implications. Section 7 concludes.

[2]Closely related, Kacperczyk et al. (2018) and Azarmsa (2022) show that information acquisition by wealthier investors can amplify inequality through higher returns. In our setting, that force is only part of the picture because the underlying costs have been trending with technology (see, among others, Frost et al. (2022) and Gambacorta et al. (2023)), Falling participation frictions pull inequality down, while falling adoption and intensity costs push it up through the between-group and within-informed channels. The net effect on the Gini, Lorenz dominance, and price informativeness depends on which margin technology moves and on the shape of the wealth distribution.

2 Baseline Economy

In this section, we develop a *parsimonious noisy–rational-expectations economy*, in the tradition of Grossman and Stiglitz (1980) and Verrecchia (1982), to isolate the channels through which falling participation and information costs shape the distribution of capital wealth and influence capital wealth inequality.

2.1 Model Assumptions

A continuum of investors, indexed by initial wealth W_{0j} and endowed with constant relative risk-aversion (CRRA) preferences, trade a risk-free bond and a single risky asset whose payoff is observed with noise. Two frictions govern their choices: (i) a *fixed stock-market entry fee* F, which determines the extensive margin of participation, and (ii) a *linear and convex information-acquisition cost* $C(x)$, which determines the extensive and intensive margins of private signal precision x about the risky asset.

Because investors differ only in wealth, the model delivers closed-form thresholds for (i) the minimum wealth required to enter the market and (ii) the higher wealth level above which information production becomes optimal. These thresholds form the basis for comparative-static exercises in which reductions in F or in the parameters of $C(\cdot)$ represent improvements in trading platforms or analytical technology. The analysis traces the resulting changes in participation, price informativeness, and the implied capital-wealth inequality, providing a transparent benchmark for the richer robustness exercises and extensions presented later in the paper.

Agents and Preferences A continuum of investors $j \in [0, 1]$ differ solely in their initial total endowments $W_{0j} \in (0, \bar{W}_0]$ under an initial wealth distribution denoted by V. Without loss of generality we normalize initial stock holdings to zero and let all initial wealth be held in the bond at $t = 0$; only net trades at $t = 1$ matter for payoffs in our one-period setting. Each investor chooses a stock position q_{1j} to maximize a constant-relative-risk-aversion (CRRA) utility, which is well approximated

by a mean–variance objective with decreasing absolute risk aversion[3]:

$$\max_{q_j} \mathbb{E}_j[W_{2j}] - \frac{\rho(W_{0j})}{2}\mathbb{V}_j[W_{2j}], \quad \text{where} \quad \rho(W_{0j}) \equiv \frac{\rho}{W_{0j}}, \quad \rho > 0. \tag{1}$$

Hence absolute risk tolerance, $1/\rho(W_{0j})$, increases with initial wealth.

Assets and Noisy Supply The financial market consists of one risk-free asset and one risky asset. The risk-free asset is in unlimited supply, with price normalized to 1 and deterministic payoff r. The risky asset has price p and a stochastic payoff z:

$$z = \mu_z + \epsilon, \text{ with } \epsilon \sim N(0, \sigma_z^2). \tag{2}$$

Noisy traders, who trade for reasons independent of payoffs or prices (e.g., for liquidity or hedging reasons), provide a stochastic supply x of the risky asset which prevents full revelation of z.

$$x = \mu_x + v, \text{ with } v \sim N(0, \sigma_x^2) \tag{3}$$

Information Technology Investor j can purchase a private signal

$$s_j = z + \delta_j, \qquad \delta_j \sim \mathcal{N}\big(0, \sigma_{s,j}^2\big),$$

where $\sigma_{s,j}^{-2} \geq 0$ is the chosen precision. Information acquisition entails the convex cost

$$C(\sigma_{s,j}^{-2}) \;=\; c_1\sigma_{s,j}^{-2} \;+\; \frac{c_0}{2}(\sigma_{s,j}^{-2})^2 \tag{4}$$

[3]This preference specification can be interpreted as a local quadratic approximation of any utility function around initial wealth, and it allows for wealth effects, similar to CRRA utility. Concretely, the heterogeneity in absolute risk aversion implies differences in the size of investors' risky portfolios, hence different gains from purchases of information (Peress, 2004, 2005). Empirically, a large body of evidence finds absolute risk aversion declines with wealth, consistent with DARA/CRRA, using portfolio, survey, and field behavior: classic household studies Friend and Blume (1975); Italian micro data relating willingness-to-pay for lotteries to wealth Guiso and Paiella (2008); administrative portfolio evidence that risky shares rise with wealth Calvet et al. (2007); and field estimates from peer-to-peer lending show absolute risk aversion falls with a wealth proxy even within investor Paravisini et al. (2017). The macro-finance wealth-distribution literature routinely adopts CRRA (a DARA class) for the same reason (e.g., Benhabib et al., 2019).

where $c_0, c_1 > 0$.[4] This implies

$$C(0) = 0, \qquad C'() \geq 0, \qquad C''() > 0 \text{ on } [0, +\infty) \text{ and } \lim_{\sigma_{s,j}^{-2} \to \infty} C'(\sigma_{s,j}^{-2}) = \infty.$$

These assumptions guarantee the existence of an interior solution and reflect the fact that more precise information is increasingly costlier.

Let $\mathcal{I}_j$ denote investor j's information set at the trading date $t = 1$.

$$\mathcal{I}_j \;=\; \begin{cases} \mathcal{I}_j(z; p) & \text{if } \sigma_{s,j}^{-2} = 0 \quad \text{(no private signal)}, \\ \mathcal{I}_j(z; p; s_j) & \text{if } \sigma_{s,j}^{-2} > 0 \quad \text{(private signal observed)}. \end{cases}$$

Thus, every market participant observes the publicly posted price p and the prior distribution of z; only those who have incurred $C(\sigma_{s,j}^{-2})$ additionally observe their own signal s_j.

Participation Cost To hold the risky asset an investor must first pay a fixed entry fee $F > 0$. If the fee is not paid, the entire portfolio is placed in the bond.

Budget Constraint and Terminal Wealth Let $\mathbf{1}_{\{\text{part}\}}$ be an indicator that equals one if investor j pays the market-entry fee F and therefore may trade the risky asset, and zero otherwise. If the investor also acquires a private signal of precision $\sigma_{s,j}^{-2}$, she incurs the information cost $C\big(\sigma_{s,j}^{-2}\big)$ defined in (4). Given the stock position q_{1j} chosen at $t = 1$, terminal wealth at $t = 2$ is

$$W_{2j} \;=\; r\Big(W_{0j} - \mathbf{1}_{\{\text{part}\}}\, F - C\big(\sigma_{s,j}^{-2}\big) - p\, q_{1j}\Big) \;+\; q_{1j}\, z, \tag{5}$$

where the first bracketed term is the payoff on the bond financed by residual wealth and the second term is the payoff on the risky asset. For a non-participant $\big(\mathbf{1}_{\{\text{part}\}} = 0,\ q_{1j} = 0,\ \sigma_{s,j}^{-2} = 0\big)$, equation (B.2) collapses to the outside option $W_{2j} = rW_{0j}$.

2.2 Timing

The sequence of events unfolds over three dates.

[4] We need $c_1 \neq 0$, which empirically guarantees a positive marginal cost of a first Bloomberg screen. Otherwise with a purely quadratic information cost, every investor ends up learning and there is no wealth threshold for information learning. With $c_1 = 0$, $W_0^{info} = 0$. This is unrealistic, although mathematically more tractable. Because we assume $c_1 > 0$, we need to find a fixed point to solve for the equilibrium.

$t = 0$ — **Planning.** Before either prices or private signals are observed, each investor decides (i) whether to incur the fixed entry fee F that grants access to the stock market and (ii) how much signal precision $\sigma_{s,j}^{-2}$ to purchase. These decisions are taken under Knightian uncertainty about the realisation of the noise-trader order flow and the asset payoff.

$t = 1$ — **Trading.** The risky-asset price p is set competitively, reflecting the aggregated information embedded in both noise-trader supply and the private signals of those investors who chose $\sigma_{s,j}^{-2} > 0$. Investors then observe the price and, if they paid the entry fee, their own private signal s_j; conditional on this information set, $\mathcal{I}_j(z; p; s_j)$, they select their stock holding q_{1j}.

$t = 2$ — **Consumption.** The fundamental payoff of the risky asset is realised, portfolios pay off, and each investor consumes terminal wealth W_{2j}.

Because planning takes place one period *before* either the price or the signal is known, equilibrium prices remain linear in the aggregate noise, preserving the tractability of the Grossman and Stiglitz (1980) framework.

2.3 Equilibrium

Equilibrium Definition An equilibrium consists of a *linear* price for the risky asset, ($rp = a + b\,z - c\,x$), where the constants a, b, and c are solved in equilibrium, such that the individual decision rules for entry (W_0^{part}), information acquisition (W_0^{info}), and portfolio choices (q_{1j}), and information precision choices ($\sigma_{s,j}^{-2}$) are mutually consistent. Here x denotes the realised noise–trader order flow. The market maker's posterior mean of z is a precision-weighted aggregator of submitted information, but z itself is the fundamental payoff.

Households. At the planning date $t = 0$ each investor anticipates how the price will translate his private information into trading profits. Given this expectation, he (i) pays the fixed fee F and enters the market if and only if initial wealth exceeds a threshold W_0^{part}; (ii) conditional on entry, purchases signal precision $\sigma_{s,j}^{-2}$ only when wealth passes a higher threshold W_0^{info}; and (iii) at the trading date $t = 1$ chooses a stock position q_{1j} that maximises mean–variance utility, exploiting the information set $\mathcal{I}_j\big(z; p\big)$ or $\mathcal{I}_j\big(z; p; s_j\big)$ according to his earlier choice.

Price formation. A competitive market maker aggregates submitted orders and sets the price so that total demand equals the exogenous noise-trader supply $x = \mu_x + v$.

Because signals and order flow are jointly normal, the unique price that clears the market is linear in (z, x) with slope $(b, -c)$, endogenously determined by investors' information choices.

Consistency. The entry cutoff W_0^{part}, the information cutoff W_0^{info}, individual demand schedules, and the price coefficients (a, b, c) are pinned down jointly:

(i) each investor's strategy maximises expected utility in (1) subject to the budget constraint in (5);

(ii) the market maker sets a market clearing stock price given by (7), such as the demand of all uninformed and informed participating investors equals the noise-traders' supply;

(iii) the implied payoffs validate the thresholds W_0^{part} and W_0^{info}.

Under the quadratic cost in (4) these conditions yield closed-form expressions for a, b, c, W_0^{part}, and W_0^{info}, producing a three-way partition of the population into non-participants, uninformed participants, and informed participants.[5]

Computation. The equilibrium is computed by solving a two–dimensional fixed point in (s, t), the aggregate risk–tolerance–weighted private precision and aggregate risk tolerance, implied by market clearing and optimal information/portfolio choices. We establish existence and local uniqueness of this fixed point in Proposition 5 (Appendix B.5), and we implement a Newton step that converges locally.

Equilibrium Solution In equilibrium, as shown in Appendix *B*.1, an investor's *optimal risky asset demand* is given by

$$q_j \;=\; \frac{W_{0j}}{\rho}\,\frac{\mathbb{E}_j[z] - rp}{\mathbb{V}_j[z]}. \tag{6}$$

In equilibrium a competitive market maker posts a *linear price* that clears the market for every realization of the fundamental payoff z and the noise–trader order flow x:

$$rp \;=\; a \;+\; b\,z \;-\; c\,x, \tag{7}$$

where the coefficients a, b, and c—together with the resulting price precision h_p—are determined exactly as in Appendix *B*.2.

[5] A two-way partition of the population into into non-participants and informed participants is also possible, provided low enough c_0, c_1.

The *optimal precision* an investor purchases, solved in Appendix B.6, is given by:

$$\sigma_{sj}^{-2}(W_0) = \frac{W_{0j}D - 2\rho c_1}{2\rho c_0}, \qquad D(s,t) \equiv \frac{\sigma_x^2 + 2st + t^2 h_0}{t^2\left(h_0 + s^2/\sigma_x^2\right)^2}, \tag{8}$$

where W_0 is initial wealth, and $D(s,t)$ represents the marginal gain per unit precision and is a function of s (the aggregate risk-tolerance weighted private precision) and t (the aggregate risk tolerance) as

$$t = \int_j \rho^{-1}(W_{0j})dV, \qquad s = \int_j \rho^{-1}(W_{0j})\sigma_{sj}^{-2}(W_0)dV. \tag{9}$$

For ease, all variables are defined in Appendix A.

Wealth Cut-offs. In equilibrium the two wealth cut-offs are

$$W_0^{\text{part}} = \frac{2\rho\, F}{D(s,t)}, \qquad W_0^{\text{info}} = \frac{2\rho\, c_1}{D(s,t)}, \tag{10}$$

where F is the fixed entry fee, c_1 the linear information cost, s the aggregate private precision, and t the aggregate risk tolerance. $D(s,t)$ represents the marginal gain per unit precision, and $h_0 > 0$ the prior precision, and $\sigma_x^2 > 0$ the variance of noise trading.[6] All variables (e.g., s, t, h_0, D, etc.) are defined in Appendix A.

Interior region and knife-edge Our comparative statics use the interior region $s > \sigma_x^2/(2t)$, under which $D_t < 0$ and $D_s < 0$ (Appendix B.4.1 Lemmas). At the knife-edge $s = \sigma_x^2/(2t)$ these slopes flatten, but all calibrations we report lie strictly inside the interior set, so Propositions 1–3 apply throughout.

2.4 Comparative Statics of Cutoffs

We study how W_0^{part} and W_0^{info} move when either the participation cost F, the extensive margin of information acquisition c_1, or the intensive margin of how much precision informed investors buy c_0, falls. A direct change in either F or c_1 cost shifts *only* the corresponding numerator, while general-equilibrium feedback effects run through s and t, which need to be solved by Newton's method as a fixed point (see Appendix B.8).

Lemma 1 (Cut in participation cost F)**.** *Holding c_1 and c_0 fixed, a decrease in F lowers W_0^{part} and raises W_0^{info}; the two cutoffs move apart.*

Proof. See Appendix B.9. A lower entry cost attracts additional (mostly uninformed) investors, increasing t and weakening the private incentives to acquire information

[6] Only the sign properties of D matter below. In the interior region $s > \sigma_x^2/(2t)$ one has $D_t < 0$ and $D_s < 0$ (Lemmas 5 and 6 in Appendix B.4.1).

(reducing s). The t channel (with $D_t < 0$) makes prices more efficient (reduces D); the s channel works in the opposite direction but is small locally. With D lower, (10) implies W_0^{part} falls directly (cheaper entry) while W_0^{info} rises (costly information is justified only at higher wealth when prices already embed more information). The reallocation is from non-participants and informed to uninformed, compressing the distribution of payoffs and reducing inequality (Gini falls).

Lemma 2 (Cut in marginal information cost c_1)**.** *Holding F and c_0 fixed, a decrease in c_1 raises W_0^{part} and lowers W_0^{info}; the two cutoffs move toward each other.*

Proof. See Appendix B.9. A lower marginal cost makes information acquisition attractive at lower wealth, so s rises. Higher s renders prices more informative (since $D_s < 0$, D declines), which reduces the benefit of entering without information; thus W_0^{part} increases. At the same time the direct effect of $c_1 \downarrow$ lowers W_0^{info}. The net reallocation is toward the informed group, widening the spread of payoffs and increasing inequality (Gini rises).

Lemma 3 (Cut in quadratic information cost c_0)**.** *Holding F and c_1 fixed, a decrease in the quadratic information cost c_0 raises both W_0^{part} and W_0^{info} proportionally.*

Proof. See Appendix B.9. Lower curvature makes precision cheaper for already-informed investors, so s increases. With $D_s < 0$, this reduces D and scales up *both* cutoffs (the adoption margin is still pinned by c_1; cf. $\frac{W_0^{\text{info}}}{W_0^{\text{part}}}$). The net effect on inequality is model-dependent: more informative prices raise the relative payoff to the informed (pushing inequality up), while higher cutoffs reduce participation at the bottom (pushing inequality down). The final effect on inequality is modest (Gini can fall or increase or remain unchanged depending on the initial wealth distribution).

Extensive vs. intensive margins The linear term c_1 sets the *marginal cost at zero precision* and therefore pins the extensive margin of information adoption. A lower c_1 reduces the wealth threshold to become informed and increases the mass of informed investors (Proposition 2). By contrast, c_0 governs the *curvature* of information costs. Holding participation fixed, a change in c_0 primarily adjusts the *intensive* precision choice of already–informed investors (and thus s), with only an indirect effect on cutoffs via $D(s,t)$.

3 Inequality: Quantitative Illustration

We now study how the three frictions in our environment—the fixed participation cost F, the marginal (per-unit) information cost c_1, and the quadratic (technology/curvature) information cost c_0—shape the cross-section of terminal wealth W_2. The initial wealth distribution W_0 is taken as given and we solve the model under Log-normal and Pareto initial wealth distributions.[7]

3.1 Quantitative Exercise

Identification and Normalization. We discipline the model with three baseline moments: (i) the mean of W_2 is normalized to 100,000 currency units; (ii) the stock-market participation rate is assumed 60%; and (iii) the share of informed investors is 5%.[8]

Given preferences, the return technology, and the outside option (all held fixed across experiments), the share targets in (ii)–(iii) pin down the endogenous wealth cutoffs for participation and information acquisition,

$$W_{\text{part}} = \frac{2\rho F}{D(s,t)} \qquad \text{and} \qquad W_{\text{info}} = \frac{2\rho c_1}{D(s,t)} \quad \Big(\Rightarrow \ \frac{W_{\text{info}}}{W_{\text{part}}} = \frac{c_1}{F} \Big),$$

where (s,t) and $D(s,t)$ are solved through a fixed point algorithm. In all counterfactuals we apply a common multiplicative rescaling to market payoffs so that $\mathbb{E}[W_2]$ remains at 100,000. This mean-preserving normalization isolates distributional effects from trivial level shifts and makes inequality comparisons across scenarios meaningful.

Counterfactuals. We implement one-at-a-time 70% cuts in F, c_1, and c_0, holding fixed the return process, signal precision, and the outside option. The F and c_1 cuts

[7] The distributions of wealth many countries, including in the US, are strikingly concentrated at the top and skewed to the right (e.g., Piketty et al. (2017), Bricker et al. (2016), and Smith et al. (2021)). For example, using the US administrative tax data, Smith et al. (2021) estimate that the top 0.1% and 1% wealth shares increased from 9.9% and 23.9% in 1989 to 15% and 31.5% in 2016, respectively.

[8] Moment (i) matches the median household net worth in the U.S. Moment (ii) is intended to match broad U.S. equity participation (direct + indirect) in household surveys and the results are qualitatively identical for different levels of equity participation (see Appendix Robustness); (iii) proxies the fraction that pays a nontrivial fixed cost to acquire professional-grade information (e.g., terminals, data feeds, or delegated active management). The Gini of W_2 is *not* separately targeted; under a lognormal prior for W_0 and our equilibrium mapping, the implied baseline is $G \simeq 0.589$. We verified numerically that forcing $G = 0.41$ while maintaining (i)–(iii) is infeasible in this parametrization—the attainable Gini set is narrow around 0.59. To reach substantially lower baseline inequality within our framework one would need a richer W_0 (e.g., a lognormal-Pareto mixture) or alternative information technology assumptions (e.g., letting c_0 enter the information cutoff).

move the participation and information margins through general equilibrium: directly and also through $D(s,t)$ as shown in Section 2.4; a c_0 cut operates only through $D(s,t)$. We explore the Gini inequality index and the Lorenz curves across these different scenarios.[9]

Proposition 1 (Inequality and cost cuts)**.** *The impact of a decrease in each of the three frictions in our environment—the fixed participation cost F, the marginal (per-unit) information cost c_1, and the quadratic (technology/curvature) information cost c_0— on the cross-section of terminal wealth W_2 is given as follows:*

1. *A decrease in F lowers W_{part} and increases W_{info}, drawing in marginal households as* uninformed *participants. This compresses between-group dispersion (non-participants vs. participants), reduces the Gini, and pushes the Lorenz curve closer to the line of equality.*

2. *A decrease in c_1 lowers W_{info} relative to W_{part}, raising the informed share among participants. This amplifies within-participants dispersion, increases the Gini, and pushes the Lorenz curve further from the line of equality.*

3. *A decrease in c_0 leaves the ratio $W_{info}/W_{part} = c_1/F$ unchanged. It acts mainly by altering (s,t) and $D(s,t)$, which tends to shift both margins in the same direction. The net effect on the Gini and the Lorenz curve is small unless the informed premium is highly sensitive to dispersion.*

Intuition. The F cut shifts mass from non-participants to (mostly) uninformed participants, compressing between-group dispersion; the c_1 cut shifts significant mass toward informed investors, increasing within-participants dispersion; the c_0 reduces both the uninformed and informed shares (the informed become better informed, however), which lowers upper-tail concentration, but raises the dispersion within the informed group. The Gini ordering follows from the comparative statics of $W_{\text{part}}(F)$ and $W_{\text{info}}(c_1)$. Below, we numerically derive the wealth distributions and inequality measures (i.e., Gini and Lorenz curves) under two different initial wealth distributions: (1) Log-normal, and (2) Pareto.

[9]The Gini index, also known as the Gini coefficient, is a statistical measure used to quantify the degree of income inequality within a population. It ranges from 0 to 1, with 0 representing perfect equality (everyone has the same income) and 1 representing perfect inequality (one person holds all the income). The index is derived from the Lorenz curve, which is a graphical representation of income or wealth inequality within a population. It plots the cumulative share of income or wealth against the cumulative share of the population, typically ordered from poorest to richest. The curve's deviation from the 45 degree line (which is the line of perfect equality where everyone has the same share) indicates the degree of inequality.

3.2 Under Lognormal Initial Wealth Distribution

Table 1 expresses the wealth distribution, the fixed-point cutoffs, and inequality across the baseline and three one-at-a-time cost cuts (all mean-preserving) under the assumption of a lognormal initial wealth distribution.

Table 1: Wealth distribution, thresholds, and inequality across scenarios (lognormal fit), scenario color code stays fixed

Scenario	s	t	D ($\times 10^{-12}$)	W_{part}	W_{info}	Share part (%)	Share info (%)	Gini
Baseline	675.18	2964.27	2.673	1403.40	9366.13	60.00	5.00	0.589
F cut	660.10	3288.21	2.641	426.09	9478.88	92.60	4.90	0.518
c_1 cut	787.13	2664.51	1.8	2084.63	4173.78	44.30	20.10	0.718
c_0 cut	820.92	2568.70	1.628	2304.94	15382.89	40.40	1.60	0.588

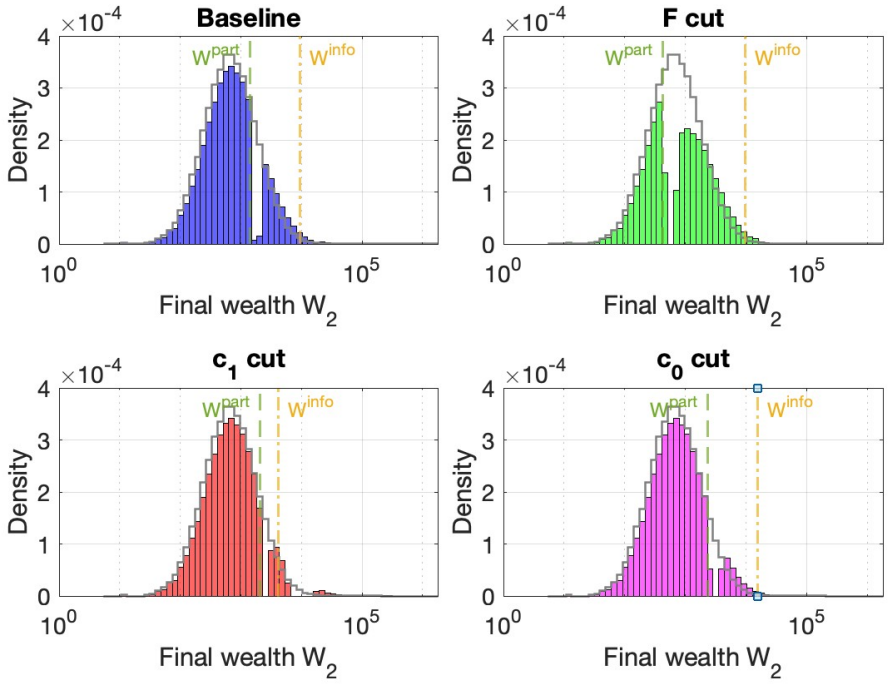
Notes. Each scenario applies a one-at-a-time 70% cut to the named cost. The magnitudes scale with the sensitivity of informed premia to dispersion. Variable definitions are in Appendix A.

Cutting the participation cost F shifts the entry threshold far left (W_{part} : $1403 \rightarrow 426$) while the information threshold barely moves (W_{info} : $9366 \rightarrow 9479$), raising participation from 60.0% to 92.6% with a slight decrease in the informed share ($5.0\% \rightarrow 4.9\%$); the Gini falls from 0.589 to 0.518, consistent with a compression of between-group gaps. Cutting the marginal information cost c_1 lowers the information threshold sharply ($9366 \rightarrow 4174$) and raises the informed share among participants (unconditional informed share $5.0\% \rightarrow 20.1\%$), even as overall participation declines ($60.0\% \rightarrow 44.3\%$) because W_{part} moves right; the right tail thickens and the Gini rises to 0.718. By contrast, cutting the quadratic information cost c_0 shifts both thresholds right (W_{part} : $1403 \rightarrow 2305$, W_{info} : $9366 \rightarrow 15383$), reducing both participation and information adoption (40.4% and 1.6%, respectively), but leaves their ratio essentially unchanged; the opposing forces on the upper tail largely offset and the Gini is virtually unchanged (0.588).

To better understand these results, in Figure 1 we plot the final wealth distributions and the endogenous cutoffs W_0^{part} and W_0^{info} superimposed across the different technology improvement scenarios. In the baseline, with W_0^{part} to the right of many households, a sizable mass remains out of the market (i.e., 40% remains out); only a small slice above W_0^{info} buys information (i.e., 5%), as in our intended calibration. The resulting W_2 is right-skewed, reflecting higher returns for market participants and especially for the informed.

A 70% F cut leads the W_0^{part} line shift left, bringing lower-wealth households into the market. Their move from the low-return outside option to the market shifts mass

Figure 1: **Final wealth distributions and endogenous cutoffs (Lognormal fit).** Each panel shows the histogram of terminal wealth W_2 (log scale) for the Baseline (blue) and for one-at-a-time 70% cuts in the participation cost F (green), the marginal information cost c_1 (red), and the quadratic information cost c_0 (magenta). Vertical lines mark the endogenous entry threshold W_{part} (green) and the information threshold W_{info} (orange). Cutting F shifts W_{part} far left while W_{info} moves little, compressing the distribution toward the center; cutting c_1 shifts W_{info} left relative to W_{part}, thickening the right tail; cutting c_0 moves both thresholds right together (their ratio is unchanged), leaving the overall shape of W_2 close to Baseline.



Notes. Baseline: lognormal initial wealth W_0 with parameters held fixed across scenarios; primitives $(\rho, h_0, \sigma_x^2, r_{\text{bond}}) = (0.98, 0.1, 1, 1.07)$ and baseline costs $(F, c_1, c_0) = (2, 2, 2)$. Scenarios apply a 70% cut to the named cost only; all other primitives are unchanged. For each case we solve the fixed-point conditions for (s, t), compute $W_{\text{part}} = 2\rho F/D(s,t)$ and $W_{\text{info}} = 2\rho c_1/D(s,t)$, map W_0 into W_2, and plot the resulting histogram using common log-spaced bins; the initial W_0 density is overlaid as a gray outline. Colors match the inequality figure (baseline blue; F cut green; c_1 cut red; c_0 cut magenta). Variable definitions are in Appendix A.

from the lower tail toward the center of the W_2 distribution. The right tail changes little, and the distribution compresses overall. This is consistent with the lower Gini: most of the action is a reduction in *between-group gaps* generated by entry.

On the other hand, a 70% cut in the marginal information cost (c_1) leads the orange W_0^{info} line to shift left, expanding the set of informed among existing participants while the participation margin moves as well. The payoffs to information push more mass into the upper tail; the left and middle of the distribution are largely unchanged. The

W_2 distribution fans out with a thicker right tail, matching the higher Gini—inequality rises through *within-participants dispersion*.

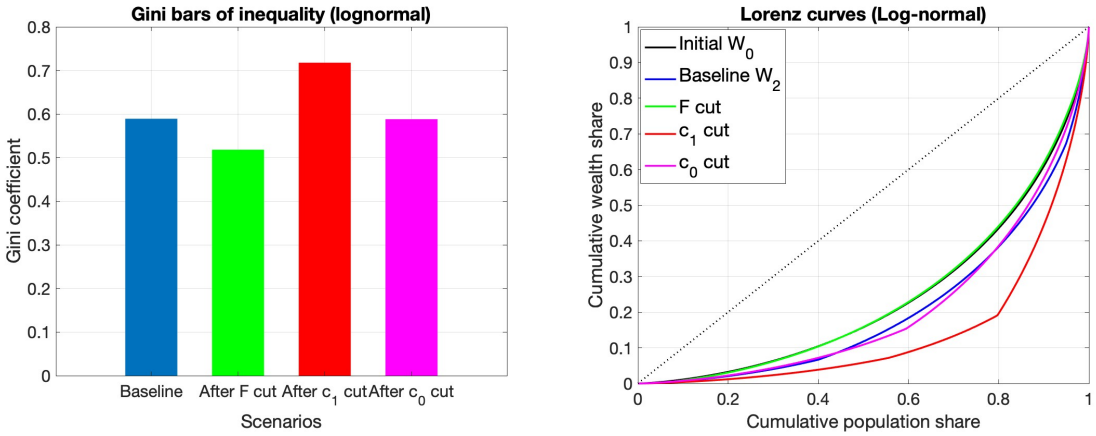
Lastly, a 70% cut in the quadratic information cost (c_0) shifts both W_0^{part} and W_0^{info} to the right by roughly the same factor. In the quantitative simulation, this shows up as lower participation and a smaller informed share, while the composition between uninformed vs. informed participants changes little on a relative basis. These shifts generate two opposing forces on the wealth distribution: (i) fewer informed investors means less mass enjoying the high information premium, which *reduces* upper-tail concentration; (ii) the fall in c_0 increases s (technology becomes more effective), which *raises* the dispersion within the informed group. These two forces nearly offset in our calibration, leaving the final wealth distribution essentially quasi-identical to the baseline.

Finally, Figure 2 summarizes how inequality (i.e., the Gini index and the Lorenz curves) responds to each cost cut. It reports the Gini coefficient (left plot) and the corresponding Lorenz curve (right plot) of W_2 in the baseline calibration and after a separate reduction in F, c_1, and c_0, holding all other parameters fixed and assuming a lognormal initial wealth distribution. In the baseline, the Gini coefficient is about 0.59. A 70% cut in the participation cost leads the Gini (in green) to fall to roughly 0.518—a decline of about -0.0710 Gini points ($\approx$ 12% relative to baseline). Lowering the fixed entry cost moves households just below the participation threshold into the market, lifting returns for the lower part of the distribution and compressing the left tail. This primarily reduces between-group inequality (non-participants vs. participants). Inequality in final wealth falls and the Lorenz curve (in green) is closer to the 45 degree line than the baseline (in blue).

On the other hand, a 70% cut in the marginal information cost (c_1) leads the Gini to rise to about 0.71—an increase of 0.129 ($\approx$ 21% relative). Cheaper information induces more already-participating – and typically wealthier – households to become informed, raising their expected returns. Gains accrue at the top, thickening the right tail (as seen in Figure 1) and increasing within-participants inequality. Distribution becomes wider, and inequality in final wealth rises. Thus, the Gini index (in red) rises, and the Lorenz curve (in red) shifts to the left of the baseline curve and away from the 45 degree line.

A comparable 70% cut in the quadratic information cost (c_0) leads the Gini at 0.588, almost unchanged from the baseline. The wealth cutoffs move to the right by roughly the same factor (the ratio $W_{\text{info}}/W_{\text{part}} = c_1/F$ is unchanged). As the two opposing forces, *reduced* upper-tail concentration and *raised* dispersion within the

Figure 2: **Inequality across scenarios.** Left: Gini coefficients of final wealth W_2 in the Baseline (blue) and after one-at-a-time 70% cuts in the participation cost F (green), the marginal information cost c_1 (red), and the quadratic information cost c_0 (magenta). Right: Lorenz curves of W_2 for the same scenarios. Cutting F compresses the distribution and lowers inequality; cutting c_1 widens it and raises inequality; cutting c_0 leaves inequality nearly unchanged. Colors are consistent across panels; the 45° line is shown for reference.



Notes. Baseline (blue): lognormal W_0; $\rho = 0.98$, $h_0 = 0.1$, $\sigma_x^2 = 1$, $(F, c_1, c_0) = (2, 2, 2)$, bond gross return $r_{\text{bond}} = 1.07$. Calibration targets: participation 60% and informed 5%. For each scenario we solve the fixed-point conditions for (s, t), compute W_{part} and W_{info}, map W_0 into W_2, and then compute the Gini and the Lorenz curve. Reported Ginis: baseline 0.589 (blue), F-cut 0.518 (green), c_1-cut 0.718 (red), c_0-cut 0.588 (magenta). Variable definitions are in Appendix A.

informed group, cancel each other the Gini (in magenta) stays essentially flat (baseline 0.589 vs. c_0 cut 0.588) and the corresponding Lorenz curve (in magenta) stays close to the baseline (in blue).

3.3 Under Pareto Initial Distribution

We now repeat the same exercise under the Pareto distribution and we compare the results under the Lognormal and Pareto initial distributions.

Table 2 expresses the wealth distribution, the fixed-point cutoffs, and inequality across the baseline and three one-at-a-time cost cuts (all mean-preserving) under the assumption of a Pareto initial wealth distribution.

Similarly, Figure 3 displays the final wealth distributions and the endogenous cutoffs W_0^{part} and W_0^{info} superimposed across the different technology improvement scenarios. Lastly, Figure 4 summarizes how inequality (i.e., the Gini index and the Lorenz curves) responds to each cost cut.

Table 2: Wealth distribution, thresholds, and inequality across scenarios (Pareto fit), scenario color code stays fixed

Scenario	s	t	D ($\times 10^{-8}$)	W_{part}	W_{info}	Share part (%)	Share info (%)	Gini
Baseline	97.05	136.30	1.72	122.67	331.44	60.00%	5.00%	0.365
F cut	90.01	185.1852	1.63	38.84	348.53	100%	4.40%	0.242
c_1 cut	119.56	91.96	1.32	159.46	159.46	31.40%	31.40%	0.583
c_0 cut	147.70	62.22	1.02	206.89	559.01	16.20%	1.40%	0.345

Notes. Each scenario applies a one-at-a-time 70% cut to the named cost. No investor is counted as informed without being a participant. Reported "Share info" can never exceed "Share part". Variable definitions are in Appendix A.

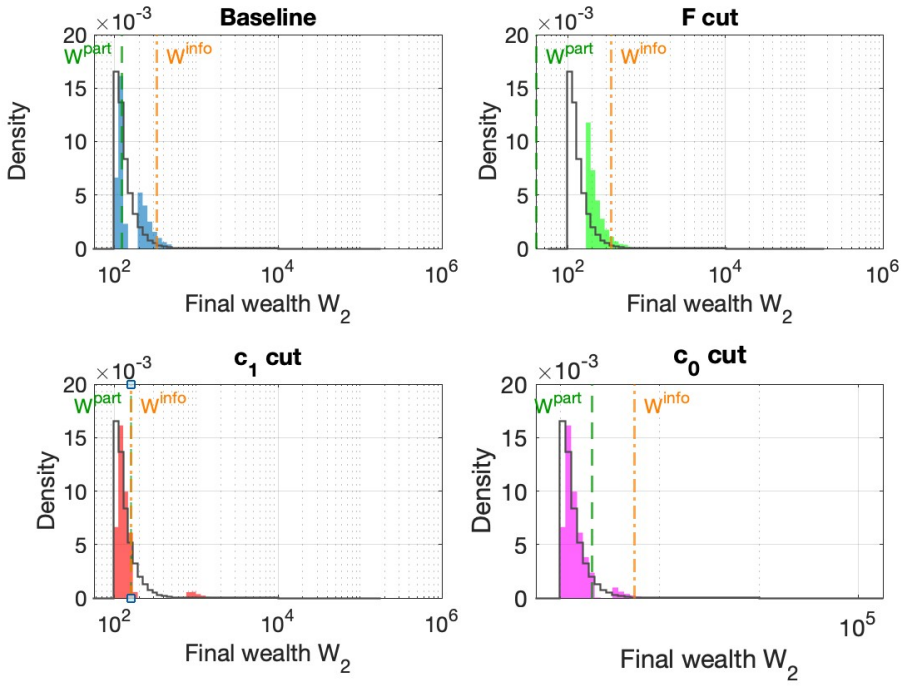
Lognormal vs. Pareto: What Changes and Why. Across priors, the policy ranking is robust: cutting the participation cost F lowers inequality, cutting the marginal information cost c_1 raises it, and cutting the quadratic information cost c_0 has at most a modest effect. The *magnitudes*, however, differ sharply because the Pareto prior places more mass in the upper tail and has compact support at the bottom.

Under the lognormal prior (Figures 1–2), the baseline Gini is 0.589. A 70% cut in F compresses between–group dispersion and lowers the Gini to 0.518 (-0.071, about -12%); a 70% cut in c_1 widens within–participants dispersion and raises the Gini to 0.718 ($+0.129$, about $+22\%$); a 70% cut in c_0 leaves the Gini essentially unchanged at 0.588.

With a Pareto prior (Figures 3–4), the baseline Gini is lower (0.365) at our shares calibration, but policy elasticities are much larger. When F falls by 70%, the participation cutoff drops below the Pareto lower bound, $W_{\min}$, so every investor enters; the Gini falls to 0.242 (-0.123, about -34%). When c_1 falls by 70%, the information cutoff plummets and (economically) almost all entrants choose to be informed; the right tail thickens markedly and the Gini jumps to 0.517 ($+0.152$, about $+42\%$). A 70% cut in c_0 shifts both cutoffs to the right and dampens the information premium; the two forces offset and the Gini moves only slightly to 0.345 (-0.020, about -5.5%).

Conceptually, the heavy tail of the Pareto prior amplifies any change that tilts returns toward the top (the c_1 cut). In contrast, the lognormal prior spreads mass more evenly across the middle, so the same policy shocks reallocate less probability into extreme tails, yielding smaller Gini responses. Finally, lowering c_0 preserves the ratio $W_{\text{info}}/W_{\text{part}} = c_1/F$, so it mostly translates into similar movements in both thresholds; that geometry leaves overall inequality nearly flat under the lognormal and only modestly lower under the Pareto calibration.

Figure 3: **Final wealth distributions and endogenous cutoffs (Pareto fit).** Each panel shows the histogram of terminal wealth W_2 (Pareto scale) for the baseline (blue) and for one-at-a-time 70% cuts in the participation cost F (green), the marginal information cost c_1 (red), and the quadratic information cost c_0 (magenta). Vertical lines mark the endogenous entry threshold W_{part} (green) and the information threshold W_{info} (orange). Cutting F shifts W_{part} far left while W_{info} moves little, compressing the distribution toward the center; cutting c_1 shifts W_{info} left relative to W_{part}, thickening the right tail; cutting c_0 moves both thresholds right together (their ratio is unchanged), leaving the overall shape of W_2 close to the baseline.



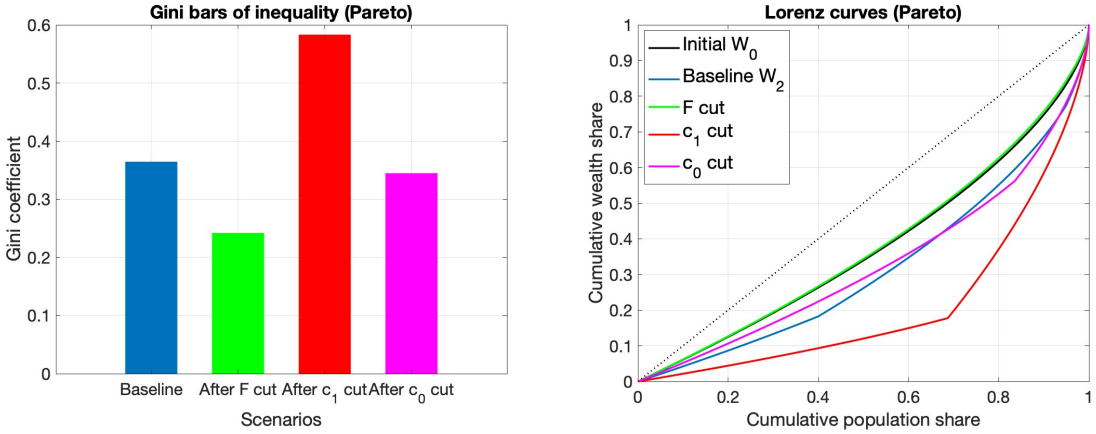
Notes. Baseline: Pareto initial wealth W_0 with parameters held fixed across scenarios; primitives $(\rho, h_0, \sigma_x^2, r_{\text{bond}}) = (0.98, 0.1, 1, 1.07)$ and baseline costs $(F, c_1, c_0) = (2, 2, 2)$. Scenarios apply a 70% cut to the named cost only; all other primitives are unchanged. For each case we solve the fixed-point conditions for (s, t), compute $W_{\text{part}} = 2\rho F/D(s,t)$ and $W_{\text{info}} = 2\rho c_1/D(s,t)$, map W_0 into W_2, and plot the resulting histogram; the initial W_0 density is overlaid as a gray outline. Colors match the inequality figure (baseline blue; F cut green; c_1 cut red; c_0 cut magenta). Variable definitions are in Appendix A.

4 Robustness, Evidence, and Generalizations

4.1 Robustness

Our results are robust to different specifications (Appendix B.14). We verify robustness of the results to (i) alternative shares of participating and/or informed investors, (ii) alternative magnitudes of cost cuts, and (iii) alternative initial wealth distributions.

Figure 4: **Inequality across scenarios (Pareto fit).** Left: Gini coefficients of final wealth W_2 in the baseline (blue) and after one-at-a-time 70% cuts in the participation cost F (green), the marginal information cost c_1 (red), and the quadratic information cost c_0 (magenta). Right: Lorenz curves of W_2 for the same scenarios. Cutting F compresses the distribution and lowers inequality; cutting c_1 widens it and raises inequality; cutting c_0 leaves inequality nearly unchanged. Colors are consistent across panels; the 45° line is shown for reference.



Notes. Baseline (blue): Pareto W_0; $\rho = 0.98$, $h_0 = 0.1$, $\sigma_x^2 = 1$, $(F, c_1, c_0) = (2, 2, 2)$, bond gross return $r_{\text{bond}} = 1.07$. Calibration targets: participation 60% and informed 5%. For each scenario we solve the fixed-point conditions for (s, t), compute W_{part} and W_{info}, map W_0 into W_2, and then compute the Gini and the Lorenz curve. Reported Ginis: baseline 0.365 (blue), F-cut 0.242 (green), c_1-cut 0.517 (red), c_0-cut 0.345 (magenta). Variable definitions are in Appendix A.

First, our robustness checks show that our main findings are not an artifact of the baseline share calibration, but a robust feature of the model's mapping from costs to the cross-section of final wealth (Appendix B.14.1). The heatmaps in Appendix Figure 15 plot the Gini coefficient under the Lognormal and Pareto initial wealth distributions when varying the shares of participating and informed investors. The results show a clear, monotone pattern under each prior: the Gini declines with higher participation and rises with a larger informed share. Iso-Gini contours slope upward, reflecting that raising participation can offset the inequality created by more information at the top. These gradients are steeper under the Pareto prior, consistent with heavier tails making both margins more powerful.[10]

Second, the Gini preserves the ranking independently of the magnitude of the cost cut which we vary between 20% and 80% for F, c_1 and c_0 (Appendix B.14.2). Appendix Table 10 shows that lowering the entry cost F compresses inequality, lowering the marginal information cost c_1 widens it, and lowering the quadratic cost c_0 does little,

[10]To isolate distributional effects, we re-normalize the mean of W_2 at each grid point; combinations that violate feasibility (i.e., informed share exceeding participation) are excluded.

across the different magnitudes of technology improvement shocks.

Third, our robustness checks show that the mechanism is not an artifact of the Lognormal prior for initial wealth (Appendix B.14.3). Replacing the Lognormal prior in Appendix Table 11 with heavy-tailed (Pareto) and thin-tailed (Uniform) benchmarks also preserves the ranking. What changes is the strength of the effects. With a heavy-tailed prior (Pareto), a reduction in c_1 shifts more mass into the right tail—amplifying the Gini increase relative to lognormal—while an F cut brings in many low-wealth households and produces a larger decline in the Gini. With a thin-tailed prior (Uniform), both tail effects are muted, so the rise in inequality from cheaper information is smaller and the decline from participation is still present but less pronounced. The policy implication is unchanged: interventions that expand market access are equality-increasing, whereas making sophisticated information cheap is equality-reducing.

4.2 Empirical Predictions and Evidence

The quantitative simulations show that, far from creating a level playing field where more readily available information simply leads to greater market efficiency, innovations in financial technology can have the opposite impact. They can create hard-to-access opportunities for long-term alpha generation for those players with the scale and resources to take advantage of it. Our framework generates sharp, testable predictions along two margins. First, reductions in participation frictions (e.g., account opening, transaction fees, digital access) should broaden stock market entry, raise participation rates particularly among lower- and middle-wealth households, and compress wealth inequality. Second, reductions in information frictions (e.g., cheaper terminals, alternative data, or sophisticated robo-advisors) disproportionately benefit wealthier or more sophisticated investors, raising their risk-adjusted returns and widening cross-sectional inequality.

Participation Evidence A growing empirical literature supports the participation margin. Commission-free trading platforms such as Robinhood boosted retail entry into equities (Ardia et al., 2025), robo-advisors expanded access to diversified portfolios (Reher and Sokolinski, 2024; D'Acunto and Rossi, 2021; D'Acunto et al., 2019; Rossi and Utkus, 2020), and cross-country evidence links digital financial inclusion to lower inequality (Suhrab et al., 2024). Descriptive evidence from the Survey of Consumer Finances (SCF) confirms that participation remains sharply skewed by wealth. Table 3 shows that fewer than 20% of households in the bottom half of the wealth distribution

hold equities, compared with over 90% of the top 10%. Participation expanded during the 1990s and 2000s, but plateaued at high levels only among the wealthy. This skew mirrors the upward shift in stock market participation documented in Melcangi and Sterk (2024), where increases in participation were disproportionately concentrated above the median income.

Table 3: Stock Market Participation by Wealth Decile (SCF, selected years)

Wealth Decile	1992	2001	2010	2022
D1 (Bottom 10%)	5%	8%	10%	14%
D2	8%	12%	15%	18%
D3	12%	18%	22%	25%
⋮				
D10 (Top 10%)	75%	82%	88%	92%

Notes. Households are sorted into wealth deciles by net worth in each Survey of Consumer Finances (SCF) wave using SCF sample weights. "Participation" equals an indicator for any equity holdings (direct stocks or indirect via mutual funds/retirement accounts). Years shown are 1992, 2001, 2010, and 2022; percentages are shares of households and are rounded to the nearest percentage point. Source: Federal Reserve Board, SCF (selected waves); authors' calculations.

Information Evidence By contrast, studies of professional adoption of alternative data and analytics point to widening return differentials. Hedge fund use of alternative data (Grennan and Michaely, 2021), satellite imagery and geolocation signals (Katona et al., 2025), and machine-learning tools in active management (Bonelli and Foucault, 2024) concentrate informational advantages and reallocate alpha toward large, data-intensive intermediaries. Evidence in Frost et al. (2022) and Gambacorta et al. (2023) also supports this channel: improvements in FinTech infrastructure change portfolio allocations and participation patterns in ways consistent with our comparative statics.

Long-term Wealth Dispersion Trends These mechanisms are consistent with broader wealth concentration dynamics. Table 4 reports the evolution of wealth shares since 1990. Data from the U.S. Census show that between 1990 and 2025Q1, wealth tripled for the top 0.1%, and doubled for the top 1%, relative to the bottom 50%. Other statistics from the Survey of Consumer Finance and from Saez and Zucman (2016) consistently show the disproportionate rise of the top 1% and top 0.1% over the past three decades.

To better capture inequality dynamics, Figure 5 presents the Gini coefficient on wealth as well as on income (left panel) and the Lorenz curves for selected years (right

Table 4: **Changes in aggregate net worth by wealth group (baseline 1990)**

Wealth group	Net Worth (1990, agg, $T)	Δ by 2000	Δ by 2010	Δ by 2019	Δ by 2025
Bottom 50%	0.7	+86%	-60%	+151%	+450%
Next 40% (50–90)	7.7	+85%	+146%	+295%	+533%
Next 9% (90–99)	7.9	+91%	+204%	+433%	+637%
Top 1% (99–99.9)	3.0	+134%	+255%	+551%	+812%
Top 0.1%	1.8	+150%	+257%	+621%	+1121%

Legend: Net worth levels are aggregate amounts for each group (trillion USD), averaged over quarters within each year; scaled to trillions for display. "Change by" columns report percent change relative to the 1990 annual average. 2025 reflects 2025:Q1 only. Wealth groups follow the DFA percentile cutoffs: Bottom 50% (p0–p50), Next 40% (p50–p90), Next 9% (p90–p99), Top 1% excl. Top 0.1% (p99–p99.9), and Top 0.1% (p99.9–p100). Aggregate net worth equals assets minus liabilities; assets include real estate, consumer durables, corporate equities and mutual fund shares, defined-benefit and defined-contribution pension entitlements, private businesses, and other assets; liabilities include home mortgages, consumer credit, and other liabilities. The first column reports the 1990 annual average in trillions of current USD (DFA series are in millions; scaled to trillions for display). This table reports *aggregate* wealth by group and is therefore not directly comparable to tables of *median* wealth by quintile. Source: Federal Reserve Board, Distributional Financial Accounts (DFA), 1989:Q3–2025:Q1; authors' calculations.

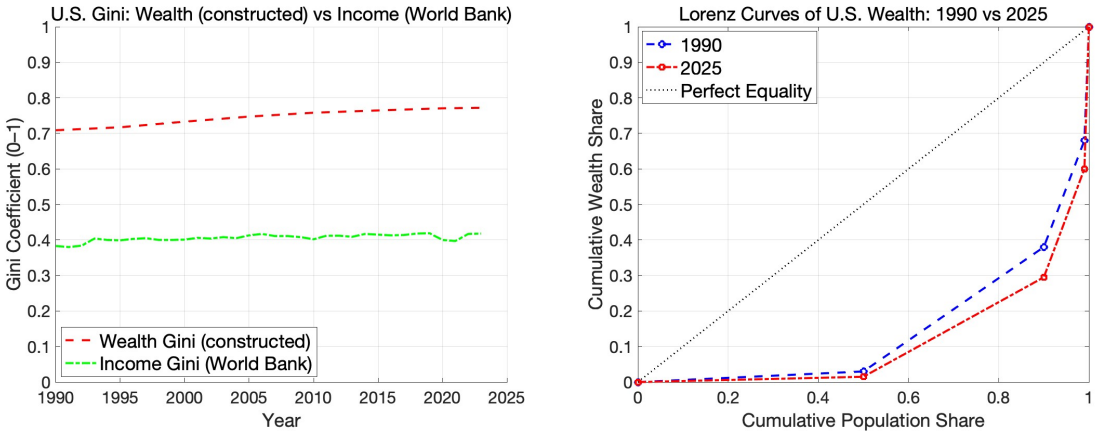
panel), showing how wealth inequality in the United States has been increasing over time.

Return Heterogeneity Finally, heterogeneous returns confirm the information-cost channel. Figure 6, based on Xavier (2021), shows that moving from the 20th to the 99th wealth percentile raises the average annual return from 3.6% to 8.3%. This 4.7 percentage-point gap reflects both portfolio composition (greater equity and private business exposure) and within-asset-class return advantages (higher yields on private business, real estate, and interest-bearing assets). These patterns corroborate our model's prediction that falling information costs disproportionately benefit the wealthy, leading to persistent divergence in cross-sectional wealth accumulation.[11] [12]

[11] See also Campbell et al. (2019) for complementary evidence from India, showing wealthy households systematically achieve higher Sharpe ratios.

[12] Theoretically, our noisy-rational expectations framework with DARA preferences also generates a thick right tail of the capital wealth distribution, consistent with U.S. data and difficult to match in canonical Bewley–Aiyagari settings. Precautionary savings models typically cannot produce wealth distributions substantially thicker than the earnings distribution they inherit (DeNardi et al., 2016; Carroll et al., 2017; Hubmer et al., 2016). By contrast, models with return heterogeneity, such as Gabaix et al. (2016) and Benhabib et al. (2017), are able to reproduce empirically relevant tails. Our setting embeds this return heterogeneity channel directly, linking it to financial information frictions and technology.

Figure 5: **Inequality trends.** Left: Gini coefficients of wealth (red) and income (green). Right: Lorenz curves of wealth in 1990 (blue) and 2025 (red). The 45 degree line represents perfect equality. Over time, wealth inequality in the United States has increased.



Notes. Income Gini (World Bank) and Wealth Gini (constructed from DFA) by year. Right: Lorenz curves for U.S. wealth in 1990 (blue) and 2025 (red); the 45° line denotes perfect equality. Summary statistics over the sample: Income Gini (World Bank): mean= 38.780, min=34.700, max=41.900. Wealth Gini (constructed): mean= 0.746, min= 0.709, max= 0.772. World Bank Gini is an income Gini; wealth inequality is typically higher. Sources: Federal Reserve Board Distributional Financial Accounts (DFA), World Bank, Saez and Zucman (2016), and author's calculations.

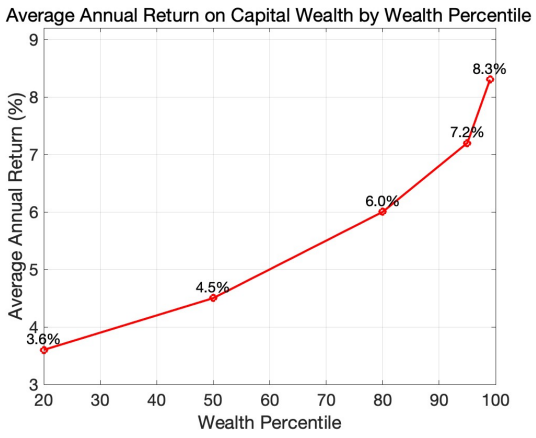
Together, these findings affirm our model's mechanisms: technological advances have expanded equity market participation, *but unevenly*. Access to data-intensive financial tools remains highly skewed toward institutions. Wealth has become increasingly concentrated over time, as shown in Lorenz curves and wealth-share trends. And wealthier individuals capture significantly higher returns, reinforcing wealth gaps. In aggregate, tech-enabled participation lowers inequality only if access expands widely; if improvements favor only the information-rich, inequality may increase, even amid greater efficiency.

4.3 Generalizations

The implications on inequality of the model baseline model are robust to a number of generalizations and to other sources of heterogeneity than wealth and risk aversion. The only requirement is that the dimensions that differentiate agents create heterogeneity in their demand for stocks. For example, differences in information costs, differences in age and lifespan, and differences in exposure to background risk all affect the demand for risky assets relative to bonds and generate similar results.

It is important to clarify differences between CARA and DARA preferences (which

Figure 6: **Average annual returns on capital wealth by wealth percentile.** Return heterogeneity is essential to explaining sustained wealth divergence.



Based on Xavier (2021). The plot shows average annual returns on capital wealth across the wealth distribution; moving from the 20th to the 99th percentile raises the average annual return from 3.6% to 8.3%. Source: Xavier (2021); authors' adaptation.

include CRRA). With CARA utility, wealth plays no role: It is irrelevant to the decisions to participate or to acquire information (at the intensive and extensive margins). However, wealth is highly relevant empirically. Lusardi et al. (2017) show that the decision to participate is significantly correlated with financial wealth. The probability of participation and the proportion of wealth invested in stocks increase with wealth, mean income, and education but decrease with the variance of income.

Thus, it is important to have a setting where financial wealth matters. CRRA preferences give relevance to wealth. But because there is no closed-form solution for equilibrium in a CRRA setting (because the price is no longer a linear function of the payoff and supply), we resort to a CARA approximation of CRRA preferences, using a strictly decreasing absolute risk aversion.

The effects of lower information costs are obtained not only in this setting, but also under CARA preferences. All that matters is the presence of a margin for participation and the coexistence of three groups of stockholders: non-stockholders, uninformed, and informed investors. Falling information costs benefit wealthier investors, who acquire more stocks, but harm direct investors who invest uninformedly, who now face a less advantageous risk-return tradeoff.

The major difference under CARA preferences is that the fixed entry cost no longer generates an information effect. The demand for information is unrelated to the expected supply of shares and to the market risk tolerance and thus to the level of participation in the stock market. This means that the variance of returns always falls

with lower entry costs.

With regards to preferences for early or late resolution of uncertainty, the results are obtained by construction because of the mean-variance preference structure. As financial information technologies improve, investors learn more about the stochastic payoffs earlier. As risk moves from the consumption to the trading period, there are larger gains from trade in earlier periods. If traders were allowed to trade before private signals were observed, then the ex-ante utility would be linear in wealth, and the traders would effectively be risk-neutral. The risk premium would equal zero. Thus, the results of this paper are obtained by construction. Lower noise (i.e., better information) does not attract investors in the stock market in this model because better information reduces the gains from trade in the second period. Moreover, better information can also reduce the risk premium in the first period.

The results extend naturally to a dynamic environment where asset returns accumulate across multiple periods. Two complementary mechanisms operate. First, on the time-series dimension, improvements in information technology shift the timing of when news is incorporated into prices. When information arrives earlier, asset prices reflect future earnings and resale values more precisely. This increases price volatility but reduces the variance of returns, because risks are resolved sooner and discounted more heavily when they are distant in time (Campbell et al. (2001)). Thus, the volatility implications of information improvements mirror those in the static model, with sharper prices but dampened return variance.

Second, on the cross-sectional dimension, repeated wealth accumulation magnifies the inequality implications of our static framework. Informed investors systematically earn higher Sharpe ratios, while uninformed investors lag behind. In a dynamic setting, these return differentials compound over time, leading to persistent divergence in wealth trajectories. Dynamic noisy rational expectations models such as Azarmsa (2022) formalize this compounding channel, showing how small per-period advantages in information acquisition translate into large long-run disparities in wealth. Empirically, heterogeneous returns across households (Xavier, 2021; Fagereng et al., 2020; Campbell et al., 2019) suggest that such mechanisms are operative in practice.

Recent evidence by Dessaint et al. (2024) further reinforces this distinction between short-run and long-run information. Short-run information technologies (e.g., high-frequency signals, real-time order flow) improve the resolution of near-term payoffs, shifting volatility forward and reducing short-run return variance—consistent with the time-series dynamics highlighted above. Long-run information technologies (e.g., deep fundamental data and machine learning applied to large histories) disproportionately

benefit resource-rich investors who can absorb the fixed costs of data processing, generating advantages that compound across periods. These long-run information advantages map directly to our cross-sectional dynamics, where inequality rises persistently as wealthy investors accumulate higher returns.

Taken together, these dynamics imply that innovations in financial information technology simultaneously reduce short-run return variance while amplifying long-run cross-sectional dispersion. While improved information stabilizes returns within a given period, it accelerates divergence across investors when the process unfolds over time, underscoring the importance of policies that sustain participation and limit extreme informational advantages.

5 Extension: Pooling into a Mutual Fund

We enrich the baseline model by allowing investors to *pool* and *delegate* their decision to a competitive informed manager who acquires a private signal on behalf of the pool. The manager charges break-even fees due to competition, as in delegated management models where fees equal costs in equilibrium (Gârleanu and Pedersen, 2018).[13]

New objects and timing. At $t = 0$, each investor j now chooses among four options: (i) stay out; (ii) enter and trade uninformed; (iii) enter and acquire precision $h_j = \sigma_{s,j}^{-2}$ individually (as in the baseline); or (iv) enter, *search* for a manager (fixed search cost $F_m \geq 0$), and delegate to a competitive fund that aggregates a mass $\mathcal{M}$ of delegators. The manager selects a pooled precision $h_m \geq 0$ and charges a per-capita fee $f \geq 0$ with zero-profit (fees cover information costs).[14] At $t = 1$ the price p is posted, the manager observes the pooled signal if $h_m > 0$, and submits the fund order; at $t = 2$ payoffs realize and fees are deducted.

Aggregation. Let $\rho^{-1}(W_{0j}) \equiv W_{0j}/\rho$ be investor j's absolute risk tolerance and $T_{\mathcal{M}} \equiv \int_{j\in\mathcal{M}} \rho^{-1}(W_{0j})\, dV(W_{0j})$ the pool's aggregate risk tolerance (sum of delegated

[13] See also Cuoco and Kaniel (2011) and He and Xiong (2013) for related delegated management environments.

[14] Equivalently, an ad-valorem fee ϕ could be used; for small ϕ the local mapping to per-capita f is standard and leaves the threshold formulas below unchanged at first order. Another alternative is a *bargaining* variant which lets the manager capture a share $\eta \in [0, 1]$ of joint surplus. The baseline corresponds to $\eta = 0$.

investors' risk tolerances). Aggregate risk tolerance and private precision become

$$t = \int_{\text{participants}} \rho^{-1}(W_{0j})\, dV(W_{0j}) \tag{11}$$

$$s = \underbrace{\int_{\text{self-informed}} \rho^{-1}(W_{0j}) h_j\, dV(W_{0j})}_{\text{individual signals}} + \underbrace{T_{\mathcal{M}} h_m}_{\text{manager's pooled signal}}. \tag{12}$$

Hence the *price rule remains* $rp = a + bz - cx$ with $b = sc$ and price precision $h_p = s^2/\sigma_x^2$ as in the baseline. In particular, the manager's activity enters s additively via $T_{\mathcal{M}} h_m$ (the price coefficients then adjust endogenously just as before).

Manager's information choice h_m. Given (s, t) and pool size $T_{\mathcal{M}}$, the manager solves the pooled analogue of the baseline information problem. Using that the marginal gain per unit precision is $D(s,t)$ (Sec. 2.3), the manager's *gross* surplus from precision h for the pool is $\frac{T_{\mathcal{M}}}{2} D(s,t)\, h$. Costs are $C(h) = c_1 h + \frac{c_0}{2} h^2$. Hence

$$\max_{h \geq 0} \frac{T_{\mathcal{M}}}{2} D(s,t)\, h - c_1 h - \frac{c_0}{2} h^2$$

$$h_m(T_{\mathcal{M}}; s, t) = \frac{T_{\mathcal{M}} D(s,t) - 2c_1}{2c_0}. \tag{13}$$

Therefore a *pool-level* information threshold emerges:

$$T_{\mathcal{M}}^{\text{info}} = \frac{2c_1}{D(s,t)}. \tag{14}$$

Pooling creates economies of scale: even if no individual j satisfies the baseline learning cutoff $(W_{0j})^{\text{info}} = \frac{2c_1\rho}{D(s,t)}$, a sufficiently large group can satisfy it and learn via the manager.

Fee setting. Competitive managers break even, so the per-capita fee equals average cost:

$$f = \frac{C(h_m)}{M} = \frac{c_1 h_m + \frac{c_0}{2}(h_m)^2}{M}, \tag{15}$$

where $M \equiv \text{mass}(\mathcal{M})$[15] and $D(s,t)$ is the marginal gain per unit precision defined in the baseline. An optional *bargaining* variant lets the manager capture a share $\eta \in [0,1]$ of joint surplus $JS = \frac{c_0}{2}(h_m)^2$, in which case $f = \frac{C(h_m)}{M} + \eta \frac{JS}{M}$. The baseline delegation model corresponds to $\eta = 0$.

[15] This "fees equal costs" condition at the industry level is the delegated-management counterpart of "alpha equals costs," a central implication in models of efficiently inefficient active management (Gârleanu and Pedersen, 2018).

Possible extension. If managers could offer a *full menu* of (h, fee) tailored to each W_{0j}, the wedge between h_m and the idiosyncratic optimum h_j would shrink and the switch region would be narrower; we keep a single h_m per fund for tractability and to match institutional practice.

Delegation decision and wealth cutoff. For investor j with tolerance $\rho^{-1}(W_{0j}) = W_{0j}/\rho$, the *incremental* mean–variance surplus from delegating (relative to uninformed) is locally $\frac{\rho^{-1}(W_{0j})}{2} D(s,t)\, h_m$, so the net surplus from delegation is

$$\Delta U_j^{\text{del}} \;=\; \frac{\rho^{-1}(W_{0j})}{2} D(s,t)\, h_m \;-\; (f + F_m).$$

Notice that the net gains of delegation over uninformed participation are linear in wealth. Thus an investor delegates if and only if

$$W_{0j} \;\geq\; W_0^{\text{del}} \;\equiv\; \rho \cdot \frac{2\,(f + F_m)}{D(s,t)\, h_m}. \tag{16}$$

Compare (16) to the baseline self-learning and participation cutoffs $W_0^{\text{info}} = \frac{2\rho c_1}{D(s,t)}$ and $W_0^{\text{part}} = \frac{2\rho F}{D(s,t)}$: delegation lowers the *per-capita* hurdle for accessing information, while search frictions (F_m) and fees (f) push in the opposite direction.

Notice also that a very rich investor can prefer to learn on their own because the manager chooses one common precision h_m to suit the pool's average risk tolerance, whereas a rich investor's idiosyncratic optimum precision growth with wealth. With quadratic costs, the self option's gain grows quadratically in wealth, while the delegation gain grown linearly. So, past a threshold, the rich switch to learning by themselves.

The net gains of self-learning over delegation for $W_{0j} \;\geq\; W_0^{\text{switch}}$ are quadratic in wealth and given by

$$\Delta U_j^{self} = \frac{[\rho^{-1}(W_{0j}) D(s,t) - 2c_1]^2}{8c_0}, \tag{17}$$

and zero otherwise.

We denote by W_{del} the *left* boundary and by W_{switch} the *right* boundary of the *delegation window*, which is the set of wealth levels $W \geq W_{\text{part}}$ such that

$$\Delta U_j^{\text{del}} \;\equiv\; \frac{W_{0j}}{\rho}\,\frac{D}{2}\, h_m \;-\; f \;-\; F_m \;\geq\; \max\{\,0,\; \Delta U_j^{self}\,\}. \tag{18}$$

Lemma 4 (Left boundary and collapse)**.** *Suppose (i) $c_1 > F$ so that $W_{info} > W_{part}$ and (ii) $T_{\mathcal{M}} D > 2c_1$ so that $h_m > 0$. Then exactly one of the following holds:*

(a) ***Interior delegation window.*** *If* $\Delta U_j^{\text{del}} > 0$*, the delegation region is nonempty and*

$$W_{del} \;=\; W_{info},$$

with a unique $W_{switch} > W_{info}$ *such that delegation strictly dominates for* $W \in [W_{info}, W_{switch})$ *and self strictly dominates for* $W > W_{switch}$*.*

(b) ***Collapse.*** *If* $\Delta U_J^{\text{del}} \leq 0$ *(or* $h_m \leq 0$*), the delegation window is empty. By convention we then report*

$$W_{del} \;=\; W_{switch},$$

and no investor delegates in equilibrium.

Proof. See Appendix D.6. The intuition is that with delegation (a), the pool's precision through the fund is strictly positive and we have a four-region partition of investors into non-participants, uninformed participants, delegating participants, and self-informed (switching) participants that prefer to self-learn over delegation. When the delegating region collapses (b), we go back to the three region partition from the baseline economy in Section 2: non-participants, uninformed participants, and self-informed participants.

5.1 Equilibrium

Equilibrium Definition An equilibrium with delegation consists of a linear price $rp = a+bz-cx$, cutoffs W_0^{part}, W_0^{info}, W_0^{del}, W_0^{switch}, aggregate objects (s,t), a manager precision h_m, a pool size M (mass of delegators), and an optimal self–precision schedule $h_j(W_{0j})$ such that:

(i) *Households.* Given (p, h_m, f), each investor chooses either not to enter, to enter uninformed, to delegate, or to self–learn. The cutoffs satisfy the delegation window in Equation 18, with $h_j(W_{0j}) = \left(\frac{W_{0j}}{\rho}D - 2c_1\right)/(2c_0)$.

(ii) *Manager.* Given the pool $\mathcal{M} = \{W \in [W_0^{\text{del}}, W_0^{\text{switch}})\}$ of mass M and aggregate tolerance $T_{\mathcal{M}}$, the fund chooses the pool's precision $h_m = (T_{\mathcal{M}}D - 2c_1)/(2c_0)$ and per–capita fee f as in Equation 15.

(iii) *Market maker.* Sets $rp = a + bz - cx$ to clear the market. With jointly normal (z, x) and linear strategies, the coefficients (a, b, c) and price precision h_p are pinned down by (s, t) exactly as in the baseline economy.

(iv) *Consistency (fixed point).* Aggregates are the sums of delegated and self-learning contributions as defined in Equations 11 and 12, and $D = D(s,t)$, and $T_{\mathcal{M}} \equiv \int_{j\in\mathcal{M}} \rho^{-1}(W_{0j})\, dV(W_{0j})$ the pool's aggregate risk tolerance (sum of delegated investors' risk tolerances); the induced cutoffs from D coincide with those used to form $\mathcal{M}$.

The quadruple (s,t,h_m,M) and $h_j(\cdot)$ (through its closed form) solve a two-by-two fixed point nested inside (s,t): the outer loop in (s,t) is as in the baseline; the inner loop pins down (h_m,M) using (13)–(15)–(16) together with the wealth distribution. Existence and local uniqueness extend from the baseline economy (Appendix B.5) by continuity. All derivations and details are in Appendix D. All formal statements for pooled delegation—manager optimal precision h_m, the pool threshold $T_{\mathcal{M}}^{\text{info}}$, competitive fees $f = C(h_m)/M$, the individual delegation cutoff W_0^{del}, and superadditivity—are collected in Appendix D.6. We use these objects below for classification and comparative statics.

5.2 Classification and Comparative Statics

Given (s,t,h_m,f), investors sort into: (i) $W_{0j} < W_0^{\text{part}}$ (non-participants); (ii) $W_0^{\text{part}} \le W_{0j} < \min\{W_0^{\text{del}}, W_0^{\text{info}}\}$ (uninformed); (iii) $W_0^{\text{del}} \le W_{0j} < W_0^{\text{switch}}$ (delegators); (iv) $W_{0j} \ge W_0^{\text{switch}}$ (self-informed). There is a unique crossing because ΔU_j^{self} is convex and steeper in W_0 beyond W_0^{info}.

Three robust predictions follow: (a) *Scale:* lower F_m or larger M raises $T_{\mathcal{M}}$, increases h_m, expands delegation; s rises and $D(s,t)$ falls. (b) *Substitution:* as s rises and D falls, W_0^{info} moves right while W_0^{del} moves left (via h_m), pushing marginal uninformed investors into delegation. (c) *Fees:* with competition, $f = C(h_m)/M$ falls in M and rises in c_1, c_0; cheaper information shifts mass from uninformed to delegated (and, if large enough, from delegated to self-informed). A formal decomposition of efficiency versus distributional effects (entry, uninformed→delegation, delegation→self, price-informativeness/risk sharing, and information costs) is provided in Appendix D.7.

Proposition 2. *In an interior equilibrium with a nonempty delegation window, the following local comparative statics hold (equivalent of Lemmas 1-3 and Proposition 1 in the baseline model):*

1. *A lower participation cost F expands entry (W_0^{part} falls); aggregate tolerance t rises; the informed share among participants changes little on impact so fees*

remain unchanged. Any feedback through $D(s,t)$ is small locally, so on net, inequality falls.

2. *A lower linear information cost c_1 boosts both self-learning and delegation, lowers D, hence decreases participation (pushes W_0^{part} right). But, h_m rises, manager fees fall, and there is more delegation at the left margin, and richer types tilt toward self–learning ($W_0^{del} = W_0^{info}$ and W_0^{switch} both fall). This ncreases payoff dispersion within participants. Inequality rises.*

3. *A lower quadratic information cost c_0 increases optimal precisions of delegating and self-informed (h_m, h_j increase). Both raises s and therefore reduces D, which reduces participation (pushes W_0^{part} right). Both delegating and self-thresholds move up, which reduces upper-tail concentration, but raises dispersion within the informed group. Inequality can grow or fall depending on whether mass is in the tails or in the middle.*

4. *A lower search cost for an informed manager F_m reduces fees, increases the pool's precision h_m; delegation expands (W_0^{del} moves left), self falls, while $D(s,t)$ falls, which makes participation fall (W_0^{part} moves right). Inequality is weakly rising.*

All proofs are in Appendix D.6.

5.3 Quantitative Exercise: Inequality with delegation

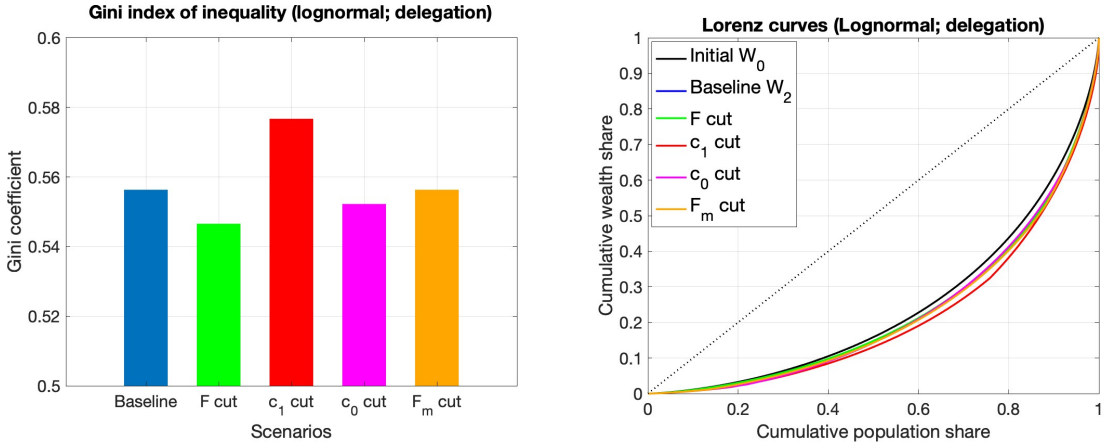
Distributional effects without ranking. Parameter changes affect inequality through three margins: (i) entry at W_0^{part} (typically compressing inequality when participation expands), (ii) moves from uninformed to delegation at W_0^{del}, and (iii) reallocation between delegation and self at W_0^{switch}. The aggregate impact on any inequality index combines the mass that moves at each cutoff (determined by the wealth density and the comparative statics for the cutoffs) with general-equilibrium feedback through (s,t), which shifts the marginal value of information $D(s,t)$. Because these masses and feedbacks are endogenous, we do not attempt to rank delegation versus self in general. The welfare-distribution decomposition used in our quantitative exercises is summarized in Appendix D.7.

5.3.1 Under Lognormal

Figure 7 displays how inequality (i.e., the Gini index and the Lorenz curves) respond to each cost cut in the model with manager delegation and pooling. Appendix Figure 17

plots the final wealth distributions and endogenous cutoffs for the Lognormal fit. The movements of the thresholds are exactly in line with the mechanisms in the model, while the *ranking* of scenarios by inequality mirrors the baseline economy without delegation.

Figure 7: **Inequality across scenarios (delegation extension).** Left: Gini coefficients of final wealth W_2 in the Baseline (blue) and after one-at-a-time 70% cuts in the participation cost F (green), and a 50% cut in the marginal information cost c_1 (red), the quadratic information cost c_0 (magenta), and the search for a manager cost F_m (orange). Right: Lorenz curves of W_2 for the same scenarios. Cutting F compresses the distribution and lowers inequality; cutting c_1 widens it and raises inequality; cutting c_0 and F_m leave inequality nearly unchanged. Colors are consistent across panels; the 45° line is shown for reference.



Notes. Baseline (blue): lognormal W_0; $\rho = 0.90$, $h_0 = 0.1$, $\sigma_x^2 = 1$, $(F, c_1, c_0) = (2, 2, 2)$, $Fm = 50$, bond gross return $r_{\text{bond}} = 1.07$. Calibration targets: participation 60% and self-informed 3.4%. For each scenario we solve the fixed-point conditions for (s, t, T_m, f, h_j, h_m), compute W_{part}, W_{del}, W_{info} and W_{switch}, map W_0 into W_2, and then compute the Gini and the Lorenz curve. Reported Ginis: baseline 0.556 (blue), F-cut 0.547 (green), c_1-cut 0.577 (red), c_0-cut 0.525 (magenta), F_m-cut 0.556 (orange). Variable definitions are in Appendix A.

Table 5 reports the four cutoffs and inequality across the five scenarios. Cutting participation costs F shifts the participation threshold $W_0^{\text{part}} = \frac{2\rho F}{D}$ sharply *left*, while the delegation window $[W_0^{\text{del}}, W_0^{\text{switch}})$ is essentially unchanged (here the left boundary is pinned at W_0^{del}=W_0^{info} by construction). Cutting the linear information cost c_1 lowers D, so W_0^{part} moves *right* (harder entry), while $W_0^{\text{del}} = \frac{2\rho c_1}{D}$ moves *left*; the switch boundary W_0^{switch} also moves *left* as richer investors' self–learning becomes relatively more attractive. Cutting the quadratic cost c_0 raises optimal precisions and lowers D, pushing *both* W_0^{part} and W_0^{del} *right*; W_0^{switch} also shifts *right*. Finally, lowering the search cost F_m for delegation affects fees but, under our implementation with W_0^{del}=W_0^{info} and a negligible self region, does not move the scalar crossing that

determines W_0^{switch}, leaving thresholds and shares numerically unchanged.

Table 5: Wealth distribution, thresholds, and inequality across scenarios (lognormal fit), scenario color code stays fixed

Scenario	s	Tm	Ts	D($x10^{-6}$)	Wp	Wdel	Wswitch	Share part (%)	Share del (%)	Share switch (%)	Gini
Baseline	9.0377	1386.17	28.2892	15.783	592.9	5907.0	53366.6	86.80%	13.60%	0%	0.556
F cut	9.0359	1386.023	28.2862	15.781	177.9	5907.7	53368.6	99%	11.90%	0%	0.547
c_1 cut	9.5461	1998.489	40.7855	12.728	735.2	3662.2	46499.4	81.60%	29.40%	0.10%	0.577
c_0 cut	9.8610	971.7403	19.8314	11.208	834.9	8318.3	60650.6	78%	8.10%	0%	0.552
F_m cut	9.0377	1386.17	28.2892	15.783	592.9	5907.0	53366.6	86.80%	13.60%	0%	0.556

Notes. Each scenario applies a one-at-a-time 50% cut to the named cost, except for the participation cost F where we apply a 70% cut. Share part is the share of participating investors (i.e., uninformed, delegated, self-learners). Share del is the share of delegating investors (i.e., pooled investors informed through a manager). Share switch is the share of self-learning investors. Share info is the sum of Share del and Share switch. Variable definitions are in Appendix A.

Crucially, the ranking of scenarios by inequality mirrors the baseline economy without delegation: reducing F delivers the strongest compression of inequality, reducing the Gini; reducing c_1 raises the Gini the most via a larger informed share and reduced participation; and changes in c_0 or F_m have much smaller effects on the Gini, with the c_0 effect theoretically ambiguous and quantitatively modest in our calibration.[16]

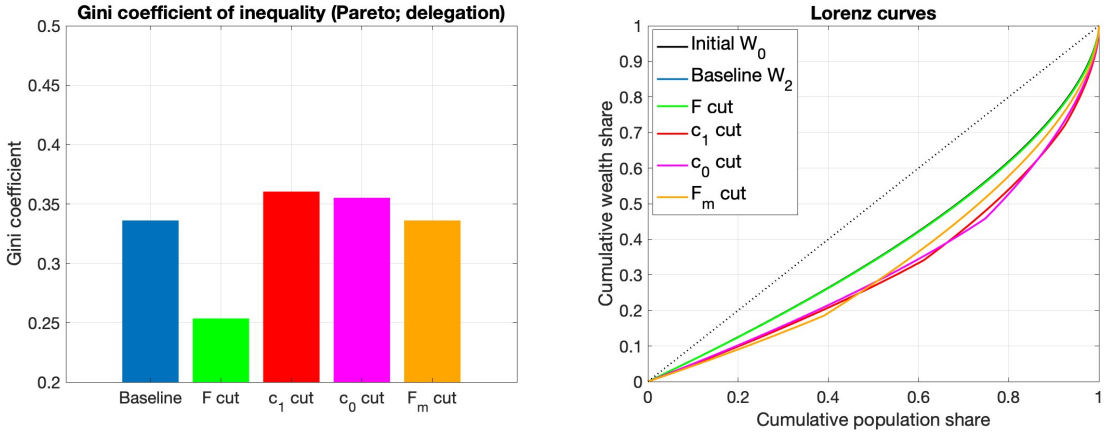
5.3.2 Under Pareto

Figure 8 plots the Gini index and the Lorenz curves in respond to each cost cut in the model with manager delegation and pooling. Similarly, Table 6 reports the four cutoffs and inequality across the five scenarios.

Table 6 shows that the ordering of scenarios is unchanged relative to the baseline model without delegation: cutting the participation cost F delivers the largest *decline* in inequality (Gini 0.336 $\rightarrow$ 0.254), while lowering information costs raises inequality, and an F_m cut is neutral in our calibration. Mechanically, an F cut shifts $W^{\text{part}} = \frac{2\rho F}{D}$ far to the left (from 122 to 37), pushing participation to 100%, with the delegation window $[W^{\text{del}}, W^{\text{switch}})$ essentially unchanged; this mass entry compresses the cross-section. By contrast, a c_1 cut increases information acquisition, raises s and lowers D: W^{part} moves right (harder entry), $W^{\text{del}} = \frac{2\rho c_1}{D}$ moves left (easier delegation), and W^{switch} falls. Participation drops to 38.8% while the delegated share rises to 19.2%, thickening the right tail and lifting the Gini to 0.36. A c_0 cut also lowers D, shifting both W^{part} and W^{del} to the right (the window moves to richer

[16] In the table, the c_0 cut yields a small decline in the Gini because the reduction in within-participants dispersion dominates the adverse-selection effect from $D \downarrow$. This is consistent with the model's ambiguous prediction for c_0.

Figure 8: **Inequality across scenarios (Pareto fit; delegation extension).** Left: Gini coefficients of final wealth W_2 in the Baseline (blue) and after one-at-a-time 70% cuts in the participation cost F (green), and a 50% cut in the marginal information cost c_1 (red), the quadratic information cost c_0 (magenta), and the search for a manager cost F_m (orange). Right: Lorenz curves of W_2 for the same scenarios. Cutting F compresses the distribution and lowers inequality; cutting c_1 widens it and raises inequality; cutting c_0 and F_m leave inequality nearly unchanged. Colors are consistent across panels; the 45° line is shown for reference.



Notes. Baseline (blue): lognormal W_0; $\rho = 0.90$, $h_0 = 0.1$, $\sigma_x^2 = 1$, $(F, c_1, c_0) = (2, 2, 2)$, $Fm = 50$, bond gross return $r_{\text{bond}} = 1.07$. Calibration targets: participation 60%. For each scenario we solve the fixed-point conditions for (s, t, T_m, f, h_j, h_m), compute W_{part}, W_{del}, W_{info} and W_{switch}, map W_0 into W_2, and then compute the Gini and the Lorenz curve. Reported Ginis: baseline 0.336 (blue), F-cut 0.254 (green), c_1-cut 0.360 (red), c_0-cut 0.355 (magenta), F_m-cut 0.336 (orange). Variable definitions are in Appendix A.

households), participation contracts to 25%, and inequality rises to 0.355. The F_m cut leaves thresholds unchanged here because the left boundary is pinned at W^{del}=W^{info} and the self-learning region is quantitatively negligible, so the Gini is unchanged.

Inequality moves through the two channels already mentioned. (i) *Between-group (entry):* when W^{part} shifts left (as with $F \downarrow$), low-wealth households enter and the distribution compresses. (ii) *Within-participants (information):* when c_1 or c_0 fall, more investors become informed and/or the informed become more precise, fattening the upper tail. Under a Pareto prior (heavy right tail), the participation contraction induced by $D \downarrow$ after $c_1 \downarrow$ or $c_0 \downarrow$ is especially powerful, so both cuts raise the Gini even with delegation.

Table 6: Wealth distribution, thresholds, and inequality across scenarios (Pareto fit), scenario color code stays fixed

Scenario	s	Tm	Ts	D($x10^{-6}$)	Wp	Wdel	Wswitch	Share part (%)	Share del (%)	Share switch (%)	Gini
Baseline	3.06	17.7339	0.36	1623.27	121.88	471.36	6397.32	61	3.4	0%	0.336
F cut	2.9981	17.4032	0.35	1603.02	37.03	477.31	6478.12	100	2	0%	0.254
c_1 cut	3.2959	38.2639	0.78	1355.24	145.99	282.29	3831.27	38.8	19.20	0%	0.360
c_0 cut	3.5772	10.3826	0.21	1136.03	174.16	673.52	9141.09	25	3.4	0%	0.355
F_m cut	3.055	17.7339	0.36	1623.27	121.88	471.36	6397.32	61	3.4	0%	0.336

Notes. Each scenario applies a one-at-a-time 50% cut to the named cost, except for the participation cost F where we apply a 70% cut. Share part is the share of participating investors (i.e., uninformed, delegated, self-learners). Share del is the share of delegating investors (i.e., pooled investors informed through a manager). Share switch is the share of self-learning investors. Share info is the sum of Share del and Share switch. Variable definitions are in Appendix A.

6 Policy Implications

The model features two margins that matter for welfare: (i) household participation in equity markets and (ii) private information acquisition in equity markets (directly or via a fund manager). Policy that lowers fixed participation costs while limiting informational advantages that disproportionately accrue to sophisticated investors raises welfare by sustaining broad participation alongside effective delegation.

Efficiency vs. welfare and inequality. Innovations that cut *participation frictions* (lower F) expand entry at the left margin and improve risk sharing; although price informativeness adjusts in general equilibrium, the welfare gains from bringing previously excluded households into the market dominate locally, and inequality compresses. By contrast, cuts in *information frictions* (lower c_1 or c_0) raise price informativeness and informational efficiency but also shift mass from uninformed to informed types and move the participation threshold rightward, so that the distributional effect—higher within-participants dispersion of certainty-equivalent payoffs—can outweigh the pure efficiency gain. Delegation tempers this force relative to self-information because its surplus gain scales roughly linearly with wealth, whereas the self-information gain is convex in wealth; consequently, for a given expansion of "being informed," self raises inequality more than delegation. A full accounting that decomposes welfare into (i) participation, (ii) uninformed→delegation, (iii) delegation→self, (iv) price-informativeness/risk-sharing, and (v) information production costs is provided in Appendix D.7, which also maps signs for small changes in F, c_1, c_0, and F_m.

Lower participation frictions. Reducing the fixed costs of entering and maintaining a brokerage relationship (e.g., digital access, account opening, and routine compliance/tax frictions) encourages participation (Abel, 2000; Vissing-Jorgensen, 2002; Bogan, 2008), and is consistent with cross-country evidence that digital financial inclusion reduces income inequality (Suhrab et al. (2024)), while recognizing that commission-free mobile brokerage expands retail participation, but needs guardrails (Ardia et al. (2025)). Policies that expand digital services and simplify retail investing administration are therefore complementary to the model's participation margin.

Build financial capability. Households with low financial literacy are markedly less likely to hold stocks (Lusardi et al., 2017), consistent with classic evidence on non-participation (Mankiw and Zeldes, 1991; Vissing-Jorgensen, 2002). Targeted financial education (curricular content, workplace programs, and just-in-time disclosures in digital channels) can raise participation and dampen volatility amplification that arises when aggregate risk is borne by a narrow investor base (Morelli, 2025).

Limit information asymmetries created by data and AI. While lower information *costs* can improve price efficiency, the model highlights that increasingly powerful data/AI tools may widen information gaps between sophisticated and retail investors. Policy should focus on (i) disclosure standards for data-driven strategies (what data, how obtained, how used), (ii) guardrails on preferential access to material, hard-to-replicate datasets, and (iii) pro-competitive access to non-proprietary public data where feasible (see, e.g., Mihet et al., 2025). The goal is not to stifle information production, but to avoid discouraging participation through large perceived informational disadvantages.

Balance efficiency and inclusion. A central implication of the model is that more information need not raise welfare if it deters participation by less sophisticated households. Competition policy heuristics that "reducing information or entry barriers cannot harm performance" need not hold once the participation margin is active (Mihet and Philippon, 2019). Regulatory cost–benefit tests for new market-data offerings or trading technologies should therefore report *both* price-efficiency gains and expected effects on the extensive margin of participation (e.g., via documented elasticities of retail entry/exit).

Product access and transparency. Existing suitability, advertising, and offering-rule frameworks that restrict complex products to sophisticated investors and require plain-language risk disclosure help align delegation with investor capability (e.g., rules in the spirit of Regulation D). When retail access to sophisticated strategies is permitted, enhanced transparency about fees, data usage, and model risks is essential.

In sum, policies that (a) lower fixed participation costs and raise financial capability, and (b) prevent extreme informational advantages from data/AI concentration, are welfare-enhancing in this setting. The paper's perspective is intentionally agnostic about industry structure: there is a single sophisticated manager, and the policy focus is on the participation–information trade-off rather than manager entry or consolidation.

7 Conclusion

We are in the midst of a historical period where advances in financial information technologies challenge virtually all aspects of business in traditional financial markets and create new markets altogether. While technology creates opportunities to dramatically reduce allocative frictions, to improve liquidity and informativeness, and to broaden households' access to financial products, it also forces society to rethink the institutions and infrastructure over which markets are built. While there is plenty of hype around financial technology, it is not clear yet whether it can make a true impact in the lives of the most financially vulnerable people. That is because financial technologies are different from other technologies. For example, financial information has special economic properties, such as nonrivalry, which imply that production possibilities are likely to be characterized by increasing returns to scale and monopoly effects, insights that have profound implications for economic growth, for capital returns, and for financial inequality.

Wealthy investors afford to acquire costly private information about asset managers and stock fundamentals. Once acquired, this private information allows them to earn higher returns, which in turn makes them wealthier, putting them in a better position to acquire even better private information. This unique property of information implies that innovations that render private information cheaper can have unintuitive externalities. They can make it easier for wealthy investors to chase and achieve high returns and pull away from the less-wealthy, who have little access to sophisticated information. My theoretical model exemplifies some possible reasons why less-wealthy investors could stop trading in the equity market because of their informational disad-

vantage. Even when they pool their information resources, they still get outcompeted by higher-wealth investors, who have more value at risk, and thus, higher incentives to acquire better private information (directly or by pooling into funds).

Future work could test the mechanisms presented in this paper using cross-sectional portfolio-level data. Scandinavian countries provide such data from tax-related forms. Unfortunately, data for US investors are difficult to obtain. But in principle, using portfolio-level data and statistics related to mobile phone and internet use, financial education, use of online banking and brokerage firm, e-trading apps, robo-advising and asset management offices, one can test whether the model's predictions for the cross-section and time series hold in the data. For example, one could test whether the wealthy have been achieving higher Sharpe ratios over time and investing an increasing share of their wealth in riskier assets, or whether poorer investors have been retrenching from risky stocks into safer assets. I base my model on aggregate trends, but the next step would be to obtain micro-level data about which, what and how investors trade.

The overall growth of investment resources and competition among investors with different wealth levels is generally considered a sign of a well-functioning financial market. This paper highlights how advances in financial information technology also have consequences beyond the democratization of capital markets, affecting the distribution of income.

References

Abel, Andrew B., "The Effects of Investing Social Security Funds in the Stock Market When Fixed Costs Prevent Some Households from Holding Stocks," Working Paper 7739, National Bureau of Economic Research June 2000.

Acemoglu, Daron, "The Simple Macroeconomics of AI," Working Paper 32487, National Bureau of Economic Research May 2024. Also published in *Economic Policy*, 2025, 40(121), 13–58.

_ , "Harms of AI," Working Paper 29247, National Bureau of Economic Research 2025. First version September 2021; revised 2025.

_ **and Pascual Restrepo**, "Robots and Jobs: Evidence from US Labor Markets," Working Paper 23285, NBER March 2017.

Admati, Anat R. and Paul Pfleiderer, "Direct and Indirect Sale of Information," *Econometrica*, 1990, *58* (4), 901–928.

Ardia, David, Clément Aymard, and Tolga Cenesizoglu, "Examining High-Frequency Patterns in Robinhood Users' Trading Behavior," 2025. SSRN working paper (first version 2022; last revised June 23, 2025).

Autor, David, "Applying AI to Rebuild Middle Class Jobs," Working Paper 32140, National Bureau of Economic Research 2024.

Azarmsa, Ehsan, "Investment Sophistication and Wealth Inequality," SSRN Working Paper No. 3206823 2022.

Badarinza, Cristian, John Y. Campbell, and Tarun Ramadorai, "International Comparative Household Finance," *Annual Review of Economics*, 2016, *8*, 111–144.

Benhabib, Jess, Alberto Bisin, and Mi Luo, "Earnings Inequality and Other Determinants of Wealth Inequality," *American Economic Review*, May 2017, *107* (5), 593–97.

_ **,** _ **, and** _ , "Wealth Distribution and Social Mobility in the US: A Quantitative Approach," *American Economic Review*, 2019, *109* (5), 1623–1647.

_ **and** _ , "Skewed Wealth Distributions: Theory and Empirics," Working Paper 21924, NBER January 2016.

Berk, Jonathan B. and Jules H. van Binsbergen, "Assessing Asset Pricing Models Using Revealed Preference," *Journal of Financial Economics*, 2016, *119* (1), 1–23.

_ **and Richard C. Green**, "Mutual Fund Flows and Performance in Rational Markets," *Journal of Political Economy*, 2004, *112* (6), 1269–1295.

Bogan, Vicki, "Stock Market Participation and the Internet," *Journal of Financial and Quantitative Analysis*, 2008, *43* (01), 191–211.

Bonelli, Maxime and Thierry Foucault, "Displaced by Big Data: Evidence from Active Fund Managers," Working Paper / SSRN No. 4527672 2024.

Bricker, Jesse, Alice Henriques, Jacob Krimmel, and John Sabelhaus, "Estimating Top Income and Wealth Shares: Sensitivity to Data and Methods," *American Economic Review*, May 2016, *106* (5), 641–45.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini, "Down or Out: Assessing the Welfare Costs of Household Investment Mistakes," *Journal of Political Economy*, 2007, *115* (5), 707–747.

Campbell, John, Martin Lettau, Burton G. Malkiel, and Yexiao Xu, "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk," *The Journal of Finance*, 2001, *56* (1), 1–43.

Campbell, John Y., Tarun Ramadorai, and Benjamin Ranish, "Do the Rich Get Richer in the Stock Market? Evidence from India," *American Economic Review: Insights*, 2019, *1* (2), 225–240.

Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White, "The distribution of wealth and the marginal propensity to consume," *Quantitative Economics*, 2017, *8* (3), 977–1020.

Cuoco, Domenico and Ron Kaniel, "Equilibrium Prices in the Presence of Delegated Portfolio Management," *Journal of Financial Economics*, 2011, *101* (2), 264–296.

D'Acunto, Francesco and Alberto G. Rossi, "Robo-Advising," in Raghavendra Rau, Robert Wardrop, and Luigi Zingales, eds., *The Palgrave Handbook of Technological Finance*, Cham: Palgrave Macmillan, 2021, pp. 725–749.

_ , Nagpurnanand Prabhala, and Alberto G. Rossi, "The Promises and Pitfalls of Robo-Advising," *Review of Financial Studies*, 2019, *32* (5), 1983–2020.

DeNardi, Maria Cristina, Giulio Fella, and Gonzalo Paz Pardo, "The Implications of Richer Earnings Dynamics for Consumption and Wealth," Working Paper 21917, NBER 2016.

Dessaint, Olivier, Thierry Foucault, and Laurent Frésard, "Does Alternative Data Improve Financial Forecasting? The Horizon Effect," *Journal of Finance*, 2024, *79* (3), 2237–2287.

Duffie, Darrell, Thierry Foucault, Laura Veldkamp, and Xavier Vives, *Technology and Finance: The Future of Banking*, London: CEPR Press, 2022.

Dugast, Jérôme and Thierry Foucault, "Equilibrium Data Mining and Data Abundance," *Journal of Finance*, 2025, *80* (1), 211–258.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri, "Heterogeneity and Persistence in Returns to Wealth," *Econometrica*, 2020, *88* (1), 115–170.

Foucault, Thierry, Leonardo Gambacorta, Wei Jiang, and Xavier Vives, *Barcelona 7: Artificial Intelligence in Finance*, London: CEPR Press, 2025.

Friend, Irwin and Marshall E. Blume, "The Demand for Risky Assets," *The American Economic Review*, 1975, *65* (5), 900–922.

Frost, Jon, Leonardo Gambacorta, and Romina Gambacorta, "On the nexus between wealth inequality, financial development and financial technology," *Journal of Economic Behavior Organization*, 2022, *202*, 429–451.

Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther, "Predictably Unequal? The Effects of Machine Learning on Credit Markets," *Journal of Finance*, 2022, *77* (1), 5–47.

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

Gaillard, Alexandre, Philipp Wangner, Christian Hellwig, and Nicolas Werquin, "Consumption, Wealth, and Income Inequality: A Tale of Tails," SSRN Working Paper No. 4636704 September 2023. First version: September 15, 2023; last revised: December 26, 2023.

Gambacorta, Leonardo, Romina Gambacorta, and Roxana Mihet, "FinTech, Investor Sophistication, and Financial Portfolio Choices," *The Review of Corporate Finance Studies*, 06 2023, *12* (4), 834–866.

Gârleanu, Nicolae and Lasse H. Pedersen, "Efficiently Inefficient Markets for Assets and Asset Management," *The Journal of Finance*, 2018, *73* (4).

_ **and Lasse Heje Pedersen**, "Active and Passive Investing: Understanding Samuelson's Dictum," *Review of Asset Pricing Studies*, 2022, *12* (2), 389–446.

Gomez, Matthieu, "Decomposing the Growth of Top Wealth Shares," *Econometrica*, 2023, *91* (3), 979–1024.

Greig, Fiona, Tarun Ramadorai, Alberto G. Rossi, Stephen P. Utkus, and Ansgar Walther, "Human Financial Advice in the Age of Automation," SSRN Working Paper No. 4301514 2024.

Grennan, Jillian and Roni Michaely, "FinTechs and the Market for Financial Analysis," *Journal of Financial and Quantitative Analysis*, 2021, *56* (6), 1877–1907.

Grossman, Sanford and Joseph Stiglitz, "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 1980, *70* (3), 393–408.

Guiso, Luigi and Monica Paiella, "Risk Aversion, Wealth, and Background Risk," *Journal of the European Economic Association*, 2008, *6* (6), 1109–1150.

He, Zhiguo and Wei Xiong, "Delegated Asset Management, Investment Mandates, and Asset Prices," *Journal of Monetary Economics*, 2013, *60* (5), 571–590.

Hubmer, Joachim, Per Krusell, and Anthony Smith, "The Historical Evolution of the Wealth Distribution: A Quantitative-Theoretic Investigation," NBER WP 23011, NBER 2016.

Kacperczyk, Marcin, Jaromir Nosal, and Luminita Stevens, "Investor sophistication and capital income inequality," *Journal of Monetary Economics*, 2018.

Katona, Zsolt, Marcus O. Painter, Panos N. Patatoukas, and Jean Zeng, "On the Capital Market Consequences of Big Data: Evidence from Outer Space," *Journal of Financial and Quantitative Analysis*, 2025, *60* (2), 551–579.

Korinek, Anton and Joseph E. Stiglitz, "Artificial Intelligence and Its Implications for Income Distribution and Unemployment," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, Cambridge, MA: National Bureau of Economic Research, 2018, pp. 349–390. Conference 2017; book edition published by University of Chicago Press, 2019.

Lusardi, Annamaria, Pierre-Carl Michaud, and Olivia Mitchell, "Optimal Financial Knowledge and Wealth Inequality," *Journal of Political Economy*, 2017, *125* (2), 431–477.

Mankiw, N. Gregory and Stephen Zeldes, "The consumption of stockholders and nonstockholders," *Journal of Financial Economics*, 1991, *29* (1), 97–112.

Melcangi, Davide and Vincent Sterk, "Stock Market Participation, Inequality, and Monetary Policy," *The Review of Economic Studies*, 06 2024, *92* (4), 2656–2690.

Mihet, Roxana and Thomas Philippon, "The Economics of Big Data and Artificial Intelligence," *Disruptive Innovation in Business and Finance in the Digital World (International Finance Review)*, Oct 2019, *20*, 29–43.

_ , Kumar Rishabh, and Orlando Gomes, "Is It AI or Data That Drives Market Power?," Research Paper 25-37, SFI April 2025. SFI Research Paper No. 25-37, Forthcoming at the Journal of Monetary Economics.

Morelli, Juan Martin, "Macroeconomic and Monetary Policy Implications of Limited Participation," Technical Report 2025. Revise and resubmit, *International Economic Review*.

Paravisini, Daniel, Veronica Rappoport, and Enrichetta Ravina, "Risk Aversion and Wealth: Evidence from Person-to-Person Lending Portfolios," *Management Science*, 2017, *63* (10), 3336–3351.

Peress, Joël, "Wealth, Information Acquisition, and Portfolio Choice," *Review of Financial Studies*, 2004, *17* (3), 879–914.

_ , "Information vs. Entry Costs: What Explains U.S. Stock Market Evolution?," *Journal of Financial and Quantitative Analysis*, September 2005, *40* (03), 563–594.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman, "Distributional National Accounts: Methods and Estimates for the United States*," *The Quarterly Journal of Economics*, 10 2017, *133* (2), 553–609.

Reher, Michael and Stanislav Sokolinski, "Robo advisors and access to wealth management," *Journal of Financial Economics*, 2024, *155*, 103829.

Rossi, Alberto G. and Stephen P. Utkus, "Who Benefits from Robo-advising? Evidence from Machine Learning," Working Paper 3552671, SSRN 2020.

_ **and Stephen Utkus**, "The diversification and welfare effects of robo-advising," *Journal of Financial Economics*, 2024, *157*, 103869.

Saez, Emmanuel and Gabriel Zucman, "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data *," *The Quarterly Journal of Economics*, 2016, *131* (2), 519–578.

Smith, Matthew, Owen Zidar, and Eric Zwick, "Top Wealth in America: New Estimates and Implications for Taxing the Rich," Working Paper 29374, NBER 2021.

Sockin, Michael and Mindy Z Xiaolan, "Delegated Learning and Contract Commonality in Asset Management*," *Review of Finance*, 03 2023, *27* (6), 1931–1975.

Suhrab, Muhammad, Pinglu Chen, and Atta Ullah, "Digital financial inclusion and income inequality nexus: Can technology innovation and infrastructure development help in achieving sustainable development goals?," *Technology in Society*, 2024, *76*, 102411.

Verrecchia, Robert E, "Information Acquisition in a Noisy Rational Expectations Economy," *Econometrica*, 1982, *50* (6), 1415–30.

Vissing-Jorgensen, Annette, "Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures," 2002. FED Working Paper 8884.

Xavier, Ines, "Wealth Inequality in the US: the Role of Heterogeneous Returns," Working Paper 3915439, SSRN 2021.

Online Appendix

A Baseline model: variable definitions

Table 7: Key symbols and their meaning (baseline model)

Symbol	Definition / interpretation
a, b, c	Coefficients in the linear price $p = a + bz - cx$
$C(\sigma_{sj}^{-2})$	Information (signal precision) acquisition cost for investor j; $C(\sigma_{sj}^{-2}) = c_1\sigma_{sj}^{-2} + \frac{c_0}{2}(\sigma_{sj}^{-2})^2$
c_0	Intensive margin (quadratic) information cost; same for all investors
c_1	Extensive margin (linear) information cost; same for all investors
D	Marginal gain per unit precision: $D(s,t) \equiv \dfrac{\sigma_x^2 + 2st + t^2 h_0}{t^2\left(h_0 + s^2/\sigma_x^2\right)^2}$
ΔU_j^{del}	Net of fees (f, F_m) surplus from delegating relative to remaining uninformed
ΔU_j^{self}	Net surplus from self-learning and self-investing
$\mathbb{E}_j[z]$	Posterior mean of z for an investor with information set $\mathcal{I}_j$
f	Manager per-capita investor fee
F	Fixed cost of equity market participation; same for all investors
F_m	Fixed cost of searching for informed manager; same for all investors
$Gini$	Gini coefficient of inequality; high Gini means high inequality; Gini $= \dfrac{1}{2\mu_2}\mathbb{E}[\,\lvert W_2 - W_2'\rvert\,]$
η	Share of joint surplus captured by manager; $\eta \in [0,1]$; under competitive fees $\eta = 0$
h_0	Prior precision of the fundamental payoff: $h_0 = \sigma_z^{-2}$
h_j	Private signal precision of investor j; $h_j = \sigma_{s,j}^{-2}$
h_m	Pooled private signal manager precision
h_p	Precision extracted from the price: $h_p = b^2/(c^2\sigma_x^2) = s^2/\sigma_x^2$
$\bar{h}$	Posterior precision of an *uninformed* investor: $\bar{h} = h_0 + h_p = \sigma_z^{-2} + s^2\sigma_x^{-2}$
$\mathcal{I}$	Set of self-informed participants
$\mathcal{I}_j$	Investor j's information set at trading date $t = 1$
j	Investor index $j \in [0,1]$
$\mathcal{M}$	Mass of delegators to a competitive fund
n	Reciprocal of aggregate precision: $n = s^{-1}$
p	Price of the risky asset (equity); linear $p = a + bz - cx$
q_{0j}	Investor j's initial position in the risky asset; normalized to 0 WLOG

(continued on next page)

(continued from previous page)

Symbol	Definition / interpretation
q_j	Investor j's time $t = 1$ position in the risky asset
r	Deterministic payoff of the bond
$\rho_j(W_{0j})$	Absolute risk aversion of j as a decreasing function of W_{0j}; $\rho_j(W_{0j}) = \frac{\rho}{W_{0j}}$ with $\rho > 0$
s	Aggregate (risk-tolerance–weighted) private precision: $s = \int_j h_j \, \rho^{-1}(W_{0j}) \, dV$
s_j	Private signal of investor j about the risky asset
σ_x^2	Variance of the supply of the risky asset (equity)
σ_{sj}^2	Variance of the private signal of investor j about the risky asset
σ_z^2	Variance of the risky asset (equity)
t	Aggregate risk tolerance: $t = \int_j \rho^{-1}(W_{0j}) \, dV$
T_j	Investor j's absolute risk tolerance: $T_j = W_{0j}/\rho$
$T_{\mathcal{M}}$	Pool's aggregate risk tolerance; $T_{\mathcal{M}} = \int_{participants} T j dj$
$T_{\mathcal{M}}^{info}$	Pool-level information threshold $T_{\mathcal{M}}^{info} = \frac{2c_1}{D(s,t)}$
$\mathcal{U}$	Set of uninformed participants
$\mathbb{V}^I$	Posterior variance of z for an informed investor: $\mathbb{V}^I = \left(\bar{h} + h_j\right)^{-1}$
$\mathbb{V}^U$	Posterior variance of z for an uninformed investor: $\mathbb{V}^U = \bar{h}^{-1}$
$\mathbb{V}_j[z]$	Posterior variance of z for an investor with information set $\mathcal{I}_j$
W_0^{del}	Minimum wealth threshold for delegation to manager
W_0^{info}	Minimum wealth threshold for private information acquisition about the equity market
W_0^{part}	Minimum wealth threshold for participation in the equity market
W_0^{switch}	Minimum wealth threshold for self-learning in the extension with delegation
W_{0j}	Initial wealth of investor j; $W_{0j} \in (0, \bar{W}_0]$
$\bar{W}_0$	Maximum initial wealth
W_2	Final wealth random variable for a household
W_2'	Independent copy of W_2 (same distribution, independent draw)
W_{2j}	Terminal wealth of investor j
x	Stochastic supply of the risky asset (equity)
z	Stochastic payoff of the risky asset (equity)
Share del	Share of delegating participants who pool their resources through a fund
Share info	Share of informed equity market participants; set at 5% in baseline exercises
Share part	Share of equity market participants; set at 60% in baseline exercises

(continued on next page)

(continued from previous page)

Symbol	Definition / interpretation
Share switch	Share of self-learning equity market investors in the delegating extension

B Baseline Model Derivations

B.1 Portfolio Demand

Proposition 3. *An investor's optimal risky asset demand is given by*

$$q_j \;=\; \frac{W_{0j}}{\rho}\,\frac{\mathbb{E}_j[z]-rp}{\mathbb{V}_j[z]}. \tag{B.1}$$

Proof. **Terminal wealth.** If the investor pays the market–entry fee F and, optionally, the information cost $C\big(\sigma_{s,j}^{-2}\big)$, residual funds are invested in the bond. Terminal wealth at $t=2$ is

$$W_{2j} \;=\; r\Big(W_{0j}-\mathbf{1}_{\{\text{part}\}}F-C\big(\sigma_{s,j}^{2}\big)-p\,q_j\Big)\;+\;q_j\,z \;=\; r\big(W_{0j}-\text{costs}\big)+q_j\big(z-rp\big), \tag{B.2}$$

where "costs" abbreviates $\mathbf{1}_{\{\text{part}\}}F + C(\sigma_{s,j}^{-2})$.

Mean variance objective. With absolute risk tolerance $\rho(W_{0j}) = W_{0j}/\rho$, investor j maximises

$$\max_{q_j}\; \mathbb{E}_j[W_{2j}] - \frac{\rho}{2W_{0j}}\mathbb{V}_j[W_{2j}], \tag{B.3}$$

where the information set is

$$\mathcal{I}_j = \begin{cases} \mathcal{I}_j(z;\,p), & \text{if } \sigma_{s,j}^{-2}=0,\\ \mathcal{I}_j(z;\,p;\,s_j), & \text{if } \sigma_{s,j}^{-2}>0. \end{cases}$$

First-order condition. Using (B.2),

$$\mathbb{E}_j[W_{2j}] = r\big(W_{0j}-\text{costs}\big)+q_j\big(\mathbb{E}_j[z]-rp\big), \quad \mathbb{V}_j[W_{2j}] = q_j^2\mathbb{V}_j[z].$$

Differentiating (B.3) with respect to q_j and setting the derivative to zero yields

$$\big(\mathbb{E}_j[z]-rp\big) - \frac{\rho}{W_{0j}}\,q_j\,\mathbb{V}_j[z] = 0. \tag{B.4}$$

Hence the *optimal demand* is

$$q_j \;=\; \frac{W_{0j}}{\rho}\,\frac{\mathbb{E}_j[z]-rp}{\mathbb{V}_j[z]}. \tag{B.5}$$

Posterior moments. Let $h_0 = 1/\sigma_z^{-2}$, $h_j = \sigma_{s,j}^{-2}$, and $\bar{h} = h_0 + \sum_{i\in\text{informed}} h_i$.

Because (z, x, s_j) are jointly normal,

$$\mathbb{E}_j[z] = \mu_z + \frac{b}{\bar{h}}\big(p - \mathbb{E}[p]\big) + \frac{h_j}{\bar{h}}\big(s_j - \mu_z\big), \qquad \mathbb{V}_j[z] = \frac{1}{\bar{h}}. \tag{B.6}$$

Substituting (B.6) in (B.5) gives the demand schedule that feeds market clearing and determines the coefficients a, b, c, and the aggregate precision $\bar{h}$ in the pricing rule.

□

B.2 Price Coefficients

Define t and s as the *aggregate risk-tolerance* and the *aggregate (risk-weighted) precision*

$$t = \int_j \rho^{-1}\big(W_{0j}\big)\, \mathrm{d}V(W_{0j}, q_{0j}) \text{ and } s = \int_j \sigma_{sj}^{-2}\, \rho^{-1}\big(W_{0j}\big)\, \mathrm{d}V(W_{0j}, q_{0j}). \tag{B.7}$$

Proposition 4. *The equilibrium price of the risky asset follows a linear form*

$$r\,p \;=\; a \;+\; b\,z \;-\; c\,x,$$

where the coefficients are given by

$$a = \mu_z(1-b) \;+\; c\,\mu_x \;-\; \frac{\mu_x}{t},$$

$$b = \frac{s\big(s\,\sigma_z^2 t + \sigma_x^2\big)}{t\big(s^2\sigma_z^2 + \sigma_x^2\big)},$$

$$c = \frac{s\,\sigma_z^2 t + \sigma_x^2}{t\big(s^2\sigma_z^2 + \sigma_x^2\big)}$$

where (s, t) are defined in Equation B.7, and there is the simple relation $b = sc$. Appendix A describes all the variables.

Proof. Recall the linear pricing rule[17]

$$r\,p \;=\; a \;+\; b\,z \;-\; c\,x, \tag{B.8}$$

[17] We could also express the linear price as $p = a + bz - cx$. However, we multiply the price by the gross risk-free return r for one reason only: it puts the stock's purchase price and its future payoff in the same units (date-2 consumption units). Nothing in the model hinges on this notation; it is just an accounting convenience that keeps the algebra short.

where $z \sim \mathcal{N}(\mu_z, \sigma_z^2)$ is the risky-asset payoff and $x \sim \mathcal{N}(\mu_x, \sigma_x^2)$ is the exogenous noise–trader supply.

Let[18]

$$h_0 \;=\; \sigma_z^{-2}, \qquad h_j \;=\; \sigma_{sj}^{-2}, \qquad \bar{h} \;=\; h_0 + h_p, \qquad h_p \;=\; \frac{s^2}{\sigma_x^2}$$

with the *aggregate risk-tolerance* and the *aggregate precision* defined as

$$t \;=\; \int_j \rho^{-1}\big(W_{0j}\big)\,\mathrm{d}V(W_{0j}, q_{0j}), \tag{B.9}$$

$$s \;=\; \int_j \sigma_{sj}^{-2}\,\rho^{-1}\big(W_{0j}\big)\,\mathrm{d}V(W_{0j}, q_{0j}), \tag{B.10}$$

and $n \equiv s^{-1}$ for later convenience.

Market clearing in terms of (t, s). For investor j the optimal demand derived in Section B.1 is

$$q_j \;=\; \frac{\rho^{-1}(W_{0j})}{\mathbb{V}_j[z]}\Big(\mathbb{E}_j[z] - rp\Big), \qquad \mathbb{V}_j[z] = \frac{1}{h_0 + h_j}.$$

Summing over all market participants (uninformed and informed) and noting that the idiosyncratic signal noises δ_j average out, aggregate demand can be written

$$\underbrace{\int_j q_j\,\mathrm{d}V}_{\text{aggregate demand}} \;=\; t\big(\mathbb{E}[z\,|\,p] - rp\big) \;+\; s\big(z - \mathbb{E}[z\,|\,p]\big). \tag{B.11}$$

Market clearing imposes $\int_j q_j\,\mathrm{d}V = x$.

The price reveals information about z. From (B.8), (z, x) is jointly normal with the scalar rp, so the linear projection of z on rp is

$$\mathbb{E}[z\,|\,p] \;=\; \kappa\, rp \quad \text{where} \quad \kappa = \frac{\mathbb{C}ov\big(z, rp\big)}{\mathbb{V}(rp)} = \frac{b\,\sigma_z^2}{b^2\sigma_z^2 + c^2\sigma_x^2}.$$

Define $\Delta p \equiv rp$ for notational brevity. Then $\mathbb{E}[z\,|\,p] - rp = (\kappa - 1)\,\Delta p$ and $z - \mathbb{E}[z\,|\,p] = z - \kappa\,\Delta p$.

[18] Note that $\bar{h} = h_0 + s$ is valid only in special cases (e.g. zero noise trading).

Impose market clearing. Plugging the previous expressions into (B.11) and collecting the coefficients of z and x yields two linear equations:

$$\Big[-bt + s - \frac{b^2 s\,\sigma_z^2}{b^2\sigma_z^2 + c^2\sigma_x^2} + \frac{b^2 t\,\sigma_z^2}{b^2\sigma_z^2 + c^2\sigma_x^2}\Big] z \;=\; 0, \tag{B.12}$$

$$\Big[ct + \frac{bcs\,\sigma_z^2}{b^2\sigma_z^2 + c^2\sigma_x^2} - \frac{bct\,\sigma_z^2}{b^2\sigma_z^2 + c^2\sigma_x^2}\Big] x \;=\; x. \tag{B.13}$$

Solving the system (B.12)–(B.13) gives[19]

$$b \;=\; \frac{s\big(s\,\sigma_z^2 t + \sigma_x^2\big)}{t\big(s^2\sigma_z^2 + \sigma_x^2\big)}, \tag{B.14a}$$

$$c \;=\; \frac{s\,\sigma_z^2 t + \sigma_x^2}{t\big(s^2\sigma_z^2 + \sigma_x^2\big)}, \tag{B.14b}$$

with the *simple identity* $b = s\,c$ (or $c = n\,b$).

To solve for the intercept a, we take expectations of the market–clearing identity $\int_j q_j dV = x$ and using the law of iterated expectations ($\mathbb{E}[\mathbb{E}[z \mid p]] = \mu_z$) yields

$$t\big(\mu_z - \mathbb{E}[rp]\big) = \mu_x.$$

Because $rp = a + bz - cx$, $\mathbb{E}[rp] = a + b\mu_z - c\mu_x$, and hence

$$a \;=\; \mu_z(1-b) \;+\; c\,\mu_x \;-\; \frac{\mu_x}{t}. \tag{B.15}$$

$$\text{If } \mu_x = 0 \quad\Longrightarrow\quad a = \mu_z(1-b).$$

The extra term $-\mu_x/t$ is a risk-premium adjustment: when the noise traders are net long on average ($\mu_x > 0$), risk-averse investors require a lower price (higher expected return) to hold that excess supply; the magnitude of the discount is inversely proportional to the aggregate risk tolerance t.

[19] All algebra was checked in SymPy; see the supplementary Mathematica notebook.

Price precision. Plugging in the coefficients from equations B.14*a*, B.14*b* and B.15 into the pricing rule,

$$rp \;=\; a \;+\; \underbrace{\frac{s\big(s\sigma_z^2 t+\sigma_x^2\big)}{t\big(s^2\sigma_z^2+\sigma_x^2\big)}}_{b(s,t)\;\begin{cases}\uparrow & \text{as } s\uparrow\\ \uparrow & \text{as } t\uparrow\end{cases}} z \;-\; \underbrace{\frac{s\sigma_z^2 t+\sigma_x^2}{t\big(s^2\sigma_z^2+\sigma_x^2\big)}}_{c(s,t)\;\begin{cases}\downarrow & \text{as } s\uparrow\\ \uparrow & \text{as } t\uparrow\end{cases}} x. \tag{B.16}$$

$$\frac{\partial b}{\partial s}>0,\quad \frac{\partial b}{\partial t}>0,\qquad \frac{\partial c}{\partial s}<0,\quad \frac{\partial c}{\partial t}>0,$$

Equation B.16 implies the following. First, as s increases (holding t fixed), the loading on fundamentals rises: $b(s,t)=\frac{s^2\sigma_z^2\,t}{(t+\sigma_x^2)(s^2\sigma_z^2+\sigma_x^2)}$ is strictly increasing in s. The noise loading $c(s,t)=\frac{s\sigma_z^2\,t}{(t+\sigma_x^2)(s^2\sigma_z^2+\sigma_x^2)}$ is hump–shaped in s (increasing for $s<\sigma_x/\sigma_z$ and decreasing for $s>\sigma_x/\sigma_z$). Price precision is $h_p=\frac{b^2}{c^2\sigma_x^2}=\frac{s^2}{\sigma_x^2}$ and therefore strictly increasing in s.

Second, as t increases (holding s fixed), both b and c increase monotonically, while their ratio remains $b/c=s$; price precision $h_p=s^2/\sigma_x^2$ is invariant to t. In other words, greater aggregate risk-bearing scales both price loadings without changing informativeness.

The posterior precision conveyed by the price is

$$h_p \;=\; \frac{b^2}{c^2\sigma_x^2} \;=\; \frac{s^2}{\sigma_x^2},\qquad \bar{h} \;=\; h_0+h_p \;=\; \sigma_z^{-2}+\frac{s^2}{\sigma_x^2}. \tag{B.17}$$

Crucially, the price precision h_p depends on s but not on t. Neither h_0 (the prior precision of the fundamental) nor h_p (the price precision) depends on μ_x; the mean noise-trader position affects only the intercept a via (B.15).

In equilibrium,

$$\bar{h} \;=\; h_0+h_p \;=\; \sigma_z^{-2}+s^2\sigma_x^{-2},$$

so the price becomes more informative (larger $\bar{h}$) when either the prior is tighter (smaller σ_z^2), noisy-trading is less noisy (smaller σ_x^2), or the pool of informed, risk-tolerant investors enlarges (higher s).

Equations (B.14)–(B.15) express the price coefficients (a,b,c) and the effective precisions $(h_0,h_j,\bar{h})$ *solely* in terms of the aggregate risk tolerance t, the aggregate precision s (and its inverse $n=s^{-1}$), and the exogenous variances (σ_z^2,σ_x^2). These expressions are the basis for all comparative-static and calibration exercises undertaken in the derivations that follow.

□

B.3 Indirect Utility

With terminal wealth (eq. B.2) $W_{2j} = r\big(W_{0j} - \text{costs}\big) + q_j\big(z - rp\big)$ and optimal demand (eq. B.5) $q_j = \tau_j \dfrac{\mathbb{E}_j[z] - rp}{\mathbb{V}_j[z]}$, $\tau_j \equiv \rho^{-1}(W_{0j})$, we obtain, *conditional on the trading-date information set* $\mathcal{I}_j$,

$$\mathbb{E}_j[W_{2j}] = r\big(W_{0j} - \text{costs}\big) + \tau_j \frac{(\mathbb{E}_j[z] - rp)^2}{\mathbb{V}_j[z]}, \qquad \mathbb{V}_j[W_{2j}] = \tau_j^2 \frac{(\mathbb{E}_j[z] - rp)^2}{\mathbb{V}_j[z]}.$$

Plugging these into the mean–variance objective $U_j = \mathbb{E}_j[W_{2j}] - \frac{\rho}{2W_{0j}}\mathbb{V}_j[W_{2j}]$ and using $\tau_j = W_{0j}/\rho$ gives the *conditional* indirect utility

$$U_j = r\big(W_{0j} - \text{costs}\big) + \frac{\tau_j}{2} \underbrace{\frac{(\mathbb{E}_j[z] - rp)^2}{\mathbb{V}_j[z]}}_{S_j^2 \text{ (squared Sharpe ratio)}}. \tag{B.18}$$

Taking the expectation over the random price realisation delivers the *planning-date* indirect utility

$$V_j = r\big(W_{0j} - \text{costs}\big) + \frac{\tau_j}{2}\,\mathbb{E}\big[S_j^2\big]. \tag{B.19}$$

Expected squared Sharpe ratio for an uninformed investor. With the linear price $rp = a + bz - cx$ and the notation $\tilde{z} = z - \mu_z$, $\tilde{x} = x - \mu_x$,

$$\mathbb{E}[z \mid p] = \mu_z + \kappa\,(rp - \mathbb{E}[rp]), \qquad \kappa = \frac{b\sigma_z^2}{b^2\sigma_z^2 + c^2\sigma_x^2} = \frac{s}{t\bar{h}}. \tag{B.20}$$

Hence $\mathbb{E}[z \mid p] - rp = (\mu_z - \mathbb{E}[rp]) + (\kappa - 1)\,(rp - \mathbb{E}[rp])$, where $\mu_z - \mathbb{E}[rp] = -\mu_x/t$ from market clearing. With $\mathbb{V}[rp] = b^2\sigma_z^2 + c^2\sigma_x^2$ and $\mathbb{V}[z \mid p] = \bar{h}^{-1}$ ($\bar{h} = h_0 + s^2/\sigma_x^2$),

$$\mathbb{E}\big[S_j^2\big]^U = \frac{(\mu_x/t)^2 + (\kappa - 1)^2\mathbb{V}[rp]}{\bar{h}^{-1}} = \frac{\sigma_x^2 + 2st + t^2 h_0}{t^2\bar{h}^2}. \tag{B.21}$$

B.4 Marginal value D of an infinitesimal private signal.

Let investor j purchase an additional precision $\mathrm{d}\sigma_s^{-2} > 0$. Posterior variance falls from $\bar{h}^{-1}$ to $(\bar{h} + \mathrm{d}\sigma_s^{-2})^{-1} = \bar{h}^{-1} - \mathrm{d}\sigma_s^{-2}/\bar{h}^2$, while the posterior mean adjusts by a term of order $\mathrm{d}\sigma_s^{-2}$. A first-order expansion of $\mathbb{E}[S_j^2]$ evaluated at $\sigma_s^{-2} = 0$ yields the *marginal*

gain per unit of added precision

$$\boxed{D(t,s,n) = \frac{\sigma_x^2 + 2st + t^2 h_0}{t^2\big(h_0 + s^2/\sigma_x^2\big)^2}}, \tag{B.22}$$

with $n \equiv s^{-1}$.

Thus, to a first approximation the value of $\mathrm{d}\sigma_s^{-2}$ units of private precision is $\frac{1}{2}\,\tau_j\,D\,\mathrm{d}\sigma_s^{-2} - C'(\sigma_s^{-2})\,\mathrm{d}\sigma_s^{-2}$.

B.4.1 Comparative statics of D

Write

$$N(s,t) \;=\; \sigma_x^2 + 2st + t^2 h_0, \qquad H(s) \;=\; h_0 + s^2/\sigma_x^2, \qquad D(s,t,n) \;=\; \frac{N(s,t)}{t^2 H(s)^2}, \quad n \equiv s^{-1},$$

with $\sigma_x^2, h_0 > 0$. Because $t, s > 0$ in equilibrium, the denominator is strictly positive; all signs are therefore determined by the numerators of the partial derivatives.

(i) Variation with aggregate risk-tolerance t.

Lemma 5. *D falls unambiguously when the market's aggregate risk-bearing capacity t rises. That is, ($\partial D/\partial t < 0$): a deeper market compresses expected Sharpe ratios.*

Proof. Proof:

Differentiate D with respect to t. Since H is independent of t,

$$\frac{\partial D}{\partial t} = \frac{N_t\, t^2 H^2 - N\, 2tH^2}{(t^2 H^2)^2} \;=\; \frac{t^2 N_t - 2tN}{t^4 H^2},$$

with $N_t \equiv \partial N/\partial t = 2s + 2th_0$.

Simplify the numerator.

$$\frac{\partial D}{\partial t} = \frac{t^2\big(2s + 2th_0\big) - 2t\big(\sigma_x^2 + 2st + t^2 h_0\big)}{(t^4 H^2)} \;=\; \frac{-2t\big(st + \sigma_x^2\big)}{t^4 H^2} \;<\; 0.$$

Because every factor in the fraction is positive except for the leading minus sign, $\partial D/\partial t$ is strictly negative for all admissible parameter values. □

(ii) Variation with price informativeness s.

Lemma 6. *D falls when the market's total price informativeness increases. That is, a larger pool of precise private information depresses D ($\partial D/\partial s < 0$): once prices embed more information, the incremental value of an extra signal is lower.*

Proof.

$$D(s,t) \;=\; \frac{N(s,t)}{t^2 H(s)^2}, \qquad N(s,t) = \sigma_x^2 + 2st + t^2 h_0, \qquad H(s) = h_0 + \frac{s^2}{\sigma_x^2}, \quad t, s, \sigma_x^2, h_0 > 0.$$

For simplicity, re-write the expressions above and take the partial derivatives of the building blocks:

$$N_s \;\equiv\; \frac{\partial N}{\partial s} = 2t, \qquad H_s \;\equiv\; \frac{\partial H}{\partial s} = \frac{2s}{\sigma_x^2}. \tag{B.23}$$

Because t is constant in this derivative, $\frac{\partial D}{\partial s}$, the quotient rule gives us

$$\frac{\partial D}{\partial s} = \frac{N_s\, t^2 H^2 - N\, t^2 (2 H H_s)}{(t^2 H^2)^2} \;=\; \frac{N_s H - 2 N H_s}{t^2 H^3}. \tag{B.24}$$

Insert N_s and H_s to get

$$\begin{aligned} \frac{\partial D}{\partial s} \;&=\; \frac{N_s H - 2 N H_s}{t^2 H^3} \;=\; \frac{(2t) H \;-\; 2N\left(\frac{2s}{\sigma_x^2}\right)}{t^2 H^3} \;= \\ &= \frac{2t\left(h_0 + \frac{s^2}{\sigma_x^2}\right) \;-\; \frac{4s}{\sigma_x^2}\left(\sigma_x^2 + 2st + t^2 h_0\right)}{t^2 H^3} \;=\; \frac{-2\left[\, h_0 t^2 + 2st + \sigma_x^2 \right]\left[\, s - \frac{\sigma_x^2}{2t} \right]}{t^2 H^3} \end{aligned} \tag{B.25}$$

The sign of $\partial D/\partial s$, in an equilibrium with active information acquisition is negative, as explained below.

Since $h_0 t^2 + 2st + \sigma_x^2 > 0$ and $t^2 H^3 > 0$, the **sign** of $\partial D/\partial s$ is governed by the bracket $s - \sigma_x^2/(2t)$.

$$\begin{cases} s > \frac{\sigma_x^2}{2t} \;\;\Longrightarrow\;\; \partial D/\partial s < 0, \\ s < \frac{\sigma_x^2}{2t} \;\;\Longrightarrow\;\; \partial D/\partial s > 0. \end{cases}$$

- *Equilibrium with active information acquisition.* Whenever investors do acquire information ($s > 0$), market clearing imposes $rp = a + bz - cx$ with $b = s\,c$. In such

equilibria one always has

$$2st \;\geq\; \sigma_x^2,$$

because the informed side of the market must at least offset the noise-trader variance to make information profitable. Under this empirically relevant condition $s - \sigma_x^2/(2t) \geq 0$ and the derivative in the box is $\partial D/\partial s < 0$: *more price informativeness ($s \uparrow$) lowers the marginal value of an additional private signal.*

- *Knife-edge case without enough precision.* If $s < \sigma_x^2/(2t)$ (too little aggregate precision relative to noise trading), the inequality flips and $\partial D/\partial s$ can be positive. But this is exactly the domain where the Grossman– Stiglitz equilibrium collapses: no trader would buy information at such low precision because the signal is swamped by noise.

So, in any equilibrium with active information acquisition,

$$\frac{\partial D}{\partial s} = -\frac{2\big(h_0 t^2 + 2st + \sigma_x^2\big)\left(s - \frac{\sigma_x^2}{2t}\right)}{t^2 H^3} < 0 \tag{B.26}$$

Thus, $\partial D/\partial s < 0$ holds throughout the economically admissible region in which private information is actually produced, completing the comparative-static result claimed in the main text.

□

These results establish that the marginal benefit of information diminishes as prices become more informative or as investors become collectively more risk tolerant.

B.5 Existence and local uniqueness of (s, t)

Proposition 5 (Existence and local uniqueness of (s,t))**.** *Assume $F > 0$, $c_0 > 0$, $c_1 > 0$, $\sigma_x^2 > 0$, $h_0 > 0$, and an initial wealth distribution with finite first two moments on $(0, \bar{W}_0]$. Then there exists a competitive linear rational–expectations equilibrium with a pair (s,t) solving the fixed–point system implied by market clearing and optimal information/portfolio choices. Moreover, the Jacobian of the fixed–point map is nonsingular at the solution, implying local uniqueness and Newton convergence from a neighborhood of (s,t).*

Proof. Define the aggregation map $\Phi(s,t) = (\Phi_s, \Phi_t)$ from the closed–form policies for $h_j(W_0)$ and q_{1j} together with truncated moment operators over the wealth distribution. Continuity holds on bounded sets; the interior signs $D_t < 0$ and $D_s < 0$ (see Appendix

Section B.4.1) imply a nonsingular Jacobian at the solution, ensuring local invertibility and Newton stability. □

B.6 Optimal precision σ_{sj}^{-2} under quadratic information cost

In the planning period the investor chooses precision $\sigma_{sj}^{-2} \geq 0$ to maximise

$$G(\sigma_{sj}^{-2}; W_0) = \frac{\tau(W_0)}{2} D \, \sigma_{sj}^{-2} \; - \; c_1 \sigma_{sj}^{-2} \; - \; \frac{c_0}{2} (\sigma_{sj}^{-2})^2, \qquad \tau(W_0) = \frac{W_0}{\rho},$$

where $D = D(s, t, n)$ is the *marginal gain per unit precision* derived in the previous section, $D = \dfrac{\sigma_x^2 + 2st + t^2 h_0}{t^2 (h_0 + s^2/\sigma_x^2)^2}$.

First-order condition.

$$\frac{\partial G}{\partial \sigma_{sj}^{-2}} = \frac{\tau(W_0)}{2} D \; - \; c_1 \; - \; c_0 \sigma_{sj}^{-2} \; = \; 0.$$

$$\sigma_{sj}^{-2}(W_0) = \frac{\frac{1}{2}\tau(W_0) D - c_1}{c_0} = \frac{W_0 D - 2\rho c_1}{2\rho c_0}. \tag{B.27}$$

Signal acquisition is interior ($\sigma_{sj}^{-2} > 0$) only for sufficiently wealthy investors; otherwise the optimum is the corner $\sigma_{sj}^{-2} = 0$.

B.7 Wealth thresholds

(i) Information threshold W_0^{info}. Set $\sigma_{sj}^{-2}(W_0) = 0$:

$$\frac{W_0^{\text{info}} D - 2\rho c_1}{2\rho c_0} = 0 \quad \Longrightarrow \quad \boxed{W_0^{\text{info}} = \frac{2\rho \, c_1}{D}.} \tag{B.28}$$

Investors with $W_0 \geq W_0^{\text{info}}$ acquire information ($\sigma_{sj}^{-2} > 0$); those with $W_0 < W_0^{\text{info}}$ do not.

(ii) Participation threshold W_0^{part}. Compare expected utility from entering *without* a signal ($\sigma_{sj}^{-2} = 0$) to staying out:

$$\frac{\tau(W_0)}{2} S_U^2 \; = \; F \quad \Longrightarrow \quad W_0^{\text{part}} = \frac{2\rho \, F}{S_U^2}, \tag{B.29}$$

where $S_U^2 = \dfrac{\sigma_x^2 + 2st + t^2 h_0}{t^2\left(h_0 + s^2/\sigma_x^2\right)^2}$ is the expected squared Sharpe ratio of an *uninformed* participant (Equation (B.21) in the main text).

Note that the expected squared Sharpe ratio for an uninformed investor and the marginal value D of an infinitesimal private signal are exactly the same. In other words, Equation B.21 and Equation B.22 are the same. See the proof in Appendix *C*.

$$D(s,t,n) = S_U^2(s,t) \tag{B.30}$$

Given (*B*.30), we can re-write Equation (*B*.29) as

$$\boxed{W_0^{\text{part}} = \frac{2\rho\, F}{D(s,t)}}. \tag{B.31}$$

Because $c_1 > F$ empirically, $W_0^{\text{info}} > W_0^{\text{part}}$, yielding the usual three-way partition:

$$\begin{aligned} W_0 < W_0^{\text{part}} &\implies \text{non-participant;} \\ W_0^{\text{part}} \le W_0 < W_0^{\text{info}} &\implies \text{uninformed participant;} \\ W_0 \ge W_0^{\text{info}} &\implies \text{informed participant.} \end{aligned}$$

B.8 Comparative statics with quadratic and linear information cost

B.8.1 Step 1. Implicit-function derivatives of (s,t) with quadratic and linear information cost

Let

$$\sigma_{sj}^{-2}(W;s,t) = \frac{W\, D(s,t) - 2\rho c_1}{2\rho c_0}\, \mathbf{1}_{\{W \ge W_0^{\text{info}}\}}, \qquad W_0^{\text{info}} = \frac{2\rho c_1}{D(s,t)}, \qquad W_0^{\text{part}} = \frac{2\rho F}{D(s,t)},$$

with $D(s,t) = \dfrac{\sigma_x^2 + 2st + t^2 h_0}{t^2\left(h_0 + s^2/\sigma_x^2\right)^2}$.

Define

$$\Psi_1(s,t;F,c_0,c_1) \;=\; t - \int_{W_0^{\text{part}}}^{\infty} \rho(W)\, g(W)\, dW,$$

$$\Psi_2(s,t;F,c_0,c_1) \;=\; s - \int_{W_0^{\text{info}}}^{\infty} \sigma_{sj}^{-2}(W;s,t)\, \rho(W)\, g(W)\, dW.$$

Equilibrium requires $\Psi_1 = 0,\ \Psi_2 = 0$. Write $\theta = (s,t)'$ and $X \in \{F, c_0, c_1\}$.

B.8.2 Implicit-function-theorem matrix

$$\theta_X = -\left[\partial_\theta \Psi\right]^{-1} \partial_X \Psi, \qquad \partial_\theta \Psi = \begin{pmatrix} \Psi_{1s} & \Psi_{1t} \\ \Psi_{2s} & \Psi_{2t} \end{pmatrix}, \quad \partial_X \Psi = \begin{pmatrix} \Psi_{1X} \\ \Psi_{2X} \end{pmatrix}.$$

All four entries of $\partial_\theta \Psi$ and all three $\Psi_{\bullet X}$ can be written as *one-dimensional* integrals; for example

$$\Psi_{1s} = -\, r\big(W_0^{\text{part}}\big)\, g\big(W_0^{\text{part}}\big)\, \frac{\partial W_0^{\text{part}}}{\partial s}, \quad \Psi_{2c_1} = -\, \sigma_{sj}^{-2}\big(W_0^{\text{info}}\big)\, r\big(W_0^{\text{info}}\big)\, g\big(W_0^{\text{info}}\big)\, \frac{\partial W_0^{\text{info}}}{\partial c_1}.$$

The algebraic expressions for the partial derivatives $\partial W_0^{\text{part}}/\partial\{s,t\}$ and $\partial W_0^{\text{info}}/\partial\{s,t\}$ follow directly from $D(s,t)$; all are closed form.

$$\Psi_{1F} = -\frac{2\rho}{D(s,t)}, \qquad \Psi_{2F} = 0, \quad \Psi_{1c_0} = 0, \qquad \Psi_{2c_0} = -\int_{W_0^{\text{info}}}^{\infty} \Big(-\frac{(\sigma_{sj}^{-2})^2}{2c_0^2}\Big) r\, g\, dW.$$

$$\Psi_{1c_1} = 0, \qquad \Psi_{2c_1} = -\, \sigma_{sj}^{-2}\big(W_0^{\text{info}}\big)\, r\big(W_0^{\text{info}}\big)\, g\big(W_0^{\text{info}}\big)\, \frac{2\rho}{D(s,t)}.$$

Numerically, compute the 2×2 Jacobian $\partial_\theta \Psi$, invert it once, and multiply by each column $\partial_X \Psi$.

B.8.3 Closed-form sign results

Because $\rho(W)$ and $g(W)$ are positive and $\partial W_0^{\text{part}}/\partial s > 0$, $\Psi_{1s} < 0$; similarly $\Psi_{2s} > 0$, $\Psi_{1t} < 0$, and $\Psi_{2t} < 0$. Therefore $\det\big[\partial_\theta \Psi\big] < 0$, guaranteeing a unique local equilibrium (s,t). Straight matrix multiplication then gives

$\operatorname{sgn} t_F = +,$	$\operatorname{sgn} s_F = -,$
$\operatorname{sgn} t_{c_1} = -,$	$\operatorname{sgn} s_{c_1} = +,$
$\operatorname{sgn} t_{c_0} = 0,$	$\operatorname{sgn} s_{c_0} = -.$

A cut in the entry fee raises aggregate risk tolerance ($t_F > 0$) but lowers aggregate precision ($s_F < 0$). A cut in the linear information cost does the opposite ($t_{c_1} < 0,\ s_{c_1} > 0$). Lowering the *quadratic* slope c_0 leaves t unchanged but raises every active investor's precision, so $s_{c_0} < 0$ and $t_{c_0} = 0$.

These signs are the inputs needed next, where we differentiate the cut-offs W_0^{part} and W_0^{info} and prove the opposite-direction movements stated in the main text.

Set–up. The cut–offs are

$$W_0^{\text{part}} = \frac{2\rho\, F}{D(s,t)}, \qquad W_0^{\text{info}} = \frac{2\rho\, c_1}{D(s,t)}, \qquad D(s,t) = \frac{\sigma_x^2 + 2st + t^2 h_0}{t^2\big(h_0 + s^2/\sigma_x^2\big)^2},$$

and (s,t) solve the two equilibrium equations in Step 1. Throughout we use the *active–information region* $s \geq \sigma_x^2/(2t)$ under which

$$D_s \equiv \frac{\partial D}{\partial s} < 0, \qquad D_t \equiv \frac{\partial D}{\partial t} < 0.$$

Let $X \in \{F, c_1\}$ and write $s_X \equiv \partial s/\partial X$, $t_X \equiv \partial t/\partial X$. Total derivatives are

$$\frac{dW_0^{\text{part}}}{dX} = \frac{2\rho\, \mathbf{1}_{\{X=F\}}}{D} - \frac{2\rho F}{D^2}\Big(D_s s_X + D_t t_X\Big), \qquad \frac{dW_0^{\text{info}}}{dX} = \frac{2\rho\, \mathbf{1}_{\{X=c_1\}}}{D} - \frac{2\rho c_1}{D^2}\Big(D_s s_X + D_t t_X\Big). \tag{B.32}$$

Sign inputs (from Step 1). With the linear and quadratic cost and monotone $\rho(W)$, we have

$$t_F < 0, \qquad s_F > 0; \qquad t_{c_1} > 0, \qquad s_{c_1} < 0.$$

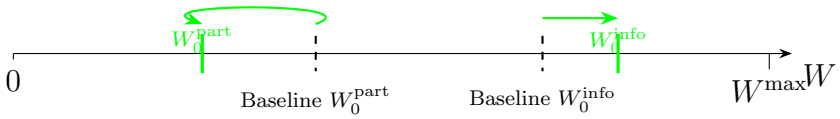
(Entry more expensive $\Rightarrow$ fewer passive participants, $t \downarrow$, but a sharper information set among incumbents, $s \uparrow$; information more expensive $\Rightarrow$ fewer informed, $s \downarrow$, and a larger passive pool, $t \uparrow$.)

B.9 Main lemmas and propositions

B.9.1 Lemma 1. Entry fee

If F decreases, then $W_0^{\text{part}} \downarrow$ and $W_0^{\text{info}} \uparrow$.

Figure 9: **Threshold movements for F cut.** Baseline thresholds W_0^{part} and W_0^{info} (dashed gray) on the wealth axis $[0, W^{\max}]$. When F falls, the wealth thresholds move *further apart* (green).



Proof. Evaluate (B.32) at $X = F$. The direct term $2\rho/D > 0$ ensures that, for an F

increase,

$$\frac{dW_0^{\text{part}}}{dF} = \frac{2\rho}{D} - \frac{2\rho F}{D^2}\Big(D_s s_F + D_t t_F\Big) \;>\; 0 \quad \text{whenever} \quad \underbrace{\big|D_s\, s_F\big|}_{\text{price effect}} \;\geq\; \underbrace{\big|D_t\, t_F\big|}_{\text{risk–bearing effect}} .$$

In the active–information region $D_s < 0$, and Step 1 gives $s_F > 0$, $t_F < 0$, so $D_s s_F < 0$ and $D_t t_F > 0$; the bracket is the *difference* of two positive magnitudes. The displayed dominance condition holds whenever the composition effect on price informativeness (new entrants are predominantly uninformed) outweighs the reduction in aggregate risk–bearing. Under that condition $dW_0^{\text{part}}/dF > 0$, so a *fall* in F yields $W_0^{\text{part}} \downarrow$.

For W_0^{info}, the derivative has no direct term:

$$\frac{dW_0^{\text{info}}}{dF} = -\frac{2\rho c_1}{D^2}\Big(D_s s_F + D_t t_F\Big).$$

With $D_s s_F < 0$ and $D_t t_F > 0$, the same dominance condition implies $dW_0^{\text{info}}/dF < 0$ fails and $dW_0^{\text{info}}/dF > 0$ holds; hence, we have $W_0^{\text{info}} \uparrow$.

This concludes the proof. □

B.9.2 Lemma 2. Linear c_1 information cost

If c_1 decreases, then $W_0^{\text{info}} \downarrow$ and $W_0^{\text{part}} \uparrow$.

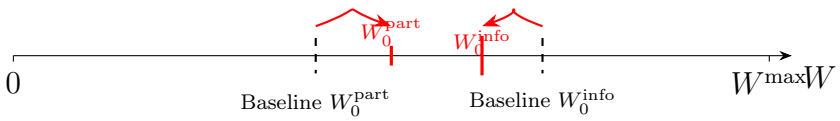


Figure 10: **Threshold movements for c_1 cut.** Baseline thresholds W_0^{part} and W_0^{info} (dashed gray) on the wealth axis $[0, W^{\max}]$. When c_1 falls, the thresholds move *closer together* (red).

Proof. Set $X = c_1$ in (B.32). For W_0^{info} we have the positive *direct* term $2\rho/D > 0$ and the feedback term with sign $-\frac{2\rho c_1}{D^2}(D_s s_{c_1} + D_t t_{c_1})$. Here $s_{c_1} < 0$ and $t_{c_1} > 0$, so $D_s s_{c_1} > 0$ and $D_t t_{c_1} < 0$. When the price–informativeness effect dominates ($|D_s s_{c_1}| \geq |D_t t_{c_1}|$), $dW_0^{\text{info}}/dc_1 > 0$, and thus a *fall* in c_1 implies $W_0^{\text{info}} \downarrow$.

For W_0^{part} there is no direct term:

$$\frac{dW_0^{\text{part}}}{dc_1} = -\frac{2\rho F}{D^2}\Big(D_s s_{c_1} + D_t t_{c_1}\Big).$$

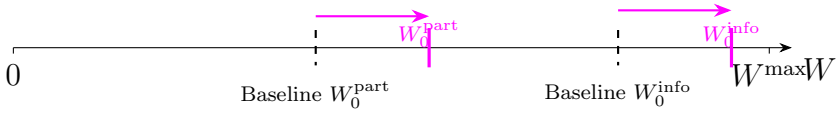


Figure 11: **Threshold movements for c_0 cut.** Baseline thresholds W_0^{part} and W_0^{info} (dashed gray) on the wealth axis $[0, W^{\max}]$. When c_0 falls, the thresholds shift *together to the right* (magenta).

With the same signs as above, the dominance condition yields $dW_0^{\text{part}}/dc_1 < 0$, hence a *fall* in c_1 implies $W_0^{\text{part}} \uparrow$.

This concludes the proof. □

B.9.3 Lemma 3. Quadratic c_0 information cost

If c_0 decreases, then W_0^{info} and W_0^{part} increase in the same direction.

Proof. Holding (F, c_1) fixed, the total derivatives satisfy

$$\frac{\partial W_0^{\text{part}}}{\partial c_0} = -\frac{2\rho F}{D(s,t)^2}\, D_{c_0}, \qquad \frac{\partial W_0^{\text{info}}}{\partial c_0} = -\frac{2\rho c_1}{D(s,t)^2}\, D_{c_0},$$

with $D_{c_0} = D_s\, s_{c_0} + D_t\, t_{c_0}$. Therefore c_0 shifts *both* cutoffs in the same direction and leaves their ratio $W_0^{\text{info}}/W_0^{\text{part}}$ unchanged. The sign of D_{c_0} is model-dependent because it aggregates the responses of s and t; without additional restrictions it is generally ambiguous.

Holding F and c_1 fixed, a decrease in c_0 operates on the intensive margin of precision. It raises s and lowers t only indirectly; consequently D moves primarily through s (higher s means lower D and lower t means higher D). Both wealth cutoffs in (10) shift upwards in the *same* direction and their ratio $\frac{W_0^{\text{info}}}{W_0^{\text{part}}} = \frac{c_1}{F}$ is unchanged. The net effect on inequality is small unless the informed premium is highly sensitive to dispersion.

For any equilibrium (s, t),

$$\frac{W_0^{\text{info}}}{W_0^{\text{part}}} = \frac{c_1}{F}, \tag{B.33}$$

independently of $D(s,t)$. Hence changes in F or c_1 *rotate* the two cutoffs relative to each other (one moves up, the other down), while changes in c_0 leave the ratio (B.33) unchanged and can only scale both cutoffs in the same direction through their common denominator $D(s,t)$.

This concludes the proof. □

Comment (sufficient conditions). The dominance conditions in the two propositions hold whenever (i) the mass that enters (exits) after an F (c_1) shock is primarily *uninformed*, and (ii) the density $g(W)$ is not too thin at the respective cut–off. Under the uniform, log–normal, and Pareto wealth distributions used in our quantitative section, the conditions are satisfied and the opposite–direction movements materialise in all calibrations we study.

Propositions B.9.1–B.9.3 deliver exactly the pattern emphasized in the main text: *after a cut in the entry fee the cut–offs move apart; after a cut in the linear information cost they move towards each other; and after a cut in the quadratic information cost, they both raise by a similar ratio.*

B.10 Equilibrium system in truncated-moment form

With absolute risk tolerance $r(W) = W/\rho$ and quadratic+linear information cost, the optimal precision for $W \geq W_0^{\text{info}}$ is $\sigma_{sj}^{-2}(W) = \big(WD - 2\rho c_1\big)/(2\rho c_0)$. Let

$$M_1(u) \equiv \int_u^\infty W\, g(W)\, dW, \qquad M_2(u) \equiv \int_u^\infty W^2\, g(W)\, dW.$$

Then the two equilibrium conditions can be written as

$$\Psi_1(s,t) = t - \frac{1}{\rho}\, M_1\big(W_0^{\text{part}}\big) = 0 \tag{B.34}$$

$$\Psi_2(s,t)s - \frac{D}{2\rho^2 c_0}\, M_2\big(W_0^{\text{info}}\big) + \frac{c_1}{c_0\rho}\, M_1\big(W_0^{\text{info}}\big) = 0, \tag{B.35}$$

where

$$W_0^{\text{part}} = \frac{2\rho F}{D(s,t)}, \qquad W_0^{\text{info}} = \frac{2\rho c_1}{D(s,t)}, \qquad D(s,t) = \frac{\sigma_x^2 + 2st + t^2 h_0}{t^2\big(h_0 + s^2/\sigma_x^2\big)^2}.$$

B.10.1 Jacobian for Newton's method

Let $\theta = (s,t)'$ and $\Psi = (\Psi_2, \Psi_1)'$. The Newton step updates $\theta \leftarrow \theta - J^{-1}\Psi$ with

$$J(\theta) = \begin{pmatrix} \Psi_{2s} & \Psi_{2t} \\ \Psi_{1s} & \Psi_{1t} \end{pmatrix}.$$

Write $u_{\mathrm{p}} = W_0^{\mathrm{part}}$ and $u_{\mathrm{i}} = W_0^{\mathrm{info}}$. Using Leibniz' rule and the fact that $\sigma_{sj}^{-2}(u_{\mathrm{i}}) = 0$,

$$\Psi_{1s} = r(u_{\mathrm{p}})g(u_{\mathrm{p}})\,\frac{\partial u_{\mathrm{p}}}{\partial s}, \qquad \Psi_{1t} = 1 + r(u_{\mathrm{p}})g(u_{\mathrm{p}})\,\frac{\partial u_{\mathrm{p}}}{\partial t},$$

$$\Psi_{2s} = 1 - \frac{D_s}{2\rho^2 c_0}\,M_2(u_{\mathrm{i}}), \qquad \Psi_{2t} = 1 - \frac{D_t}{2\rho^2 c_0}\,M_2(u_{\mathrm{i}}),$$

with

$$\frac{\partial u_{\mathrm{p}}}{\partial s} = -\,\frac{2\rho F}{D^2}\,D_s, \qquad \frac{\partial u_{\mathrm{p}}}{\partial t} = -\,\frac{2\rho F}{D^2}\,D_t,$$

$$D_s = -\frac{2\big(h_0 t^2 + 2st + \sigma_x^2\big)\,\big(s - \sigma_x^2/(2t)\big)}{t^2\big(h_0 + s^2/\sigma_x^2\big)^3}, \qquad D_t = -\frac{2\big(st + \sigma_x^2\big)}{t^3\big(h_0 + s^2/\sigma_x^2\big)^2}.$$

Algorithm (Newton with backtracking). Pick an initial guess $(s^{(0)}, t^{(0)})$ (e.g. $t^{(0)} = \mathbb{E}[W]/\rho$, $s^{(0)} = t^{(0)}\,\sigma_x/(h_0^{1/2} + \sigma_x)$), iterate

$$\theta^{(k+1)} = \theta^{(k)} - \alpha^{(k)}\,J(\theta^{(k)})^{-1}\Psi(\theta^{(k)}),$$

where $\alpha^{(k)} \in (0, 1]$ is chosen by backtracking to ensure $\|\Psi(\theta^{(k+1)})\| < \|\Psi(\theta^{(k)})\|$. Stop when $\|\Psi\|_\infty < 10^{-10}$. This produces the fixed point (s, t) for any (F, c_0, c_1) and any wealth density g.

B.10.2 Truncated moments M_1, M_2 for the three wealth priors

Log-normal $\log W \sim \mathcal{N}(\mu, \sigma^2)$**:** Let $\bar{\Phi}$ denote the upper-tail normal CDF. Then

$$M_k(u) = \exp\!\big(k\mu + \tfrac{1}{2}k^2\sigma^2\big)\,\bar{\Phi}\!\left(\frac{\ln u - \mu - k\sigma^2}{\sigma}\right), \quad k \in \{1, 2\}.$$

Pareto$(\alpha, w_{\min})$ **with upper cap** $W_{\max}$ **(to keep** M_2 **finite if** $\alpha \le 2$**):** for $u_+ \equiv \max\{u, w_{\min}\}$,

$$M_1(u) = \frac{\alpha\,w_{\min}^{\alpha}}{\alpha - 1}\,u_+^{1-\alpha} \ \text{ if } \alpha > 1,$$

$$M_2(u) = \begin{cases} \dfrac{\alpha\,w_{\min}^{\alpha}}{\alpha - 2}\,u_+^{2-\alpha}, & \alpha > 2, \\[2ex] \alpha\,w_{\min}^{\alpha}\,\ln\!\Big(\dfrac{W_{\max}}{u_+}\Big), & \alpha = 2, \\[2ex] \dfrac{\alpha\,w_{\min}^{\alpha}}{2 - \alpha}\,\Big(W_{\max}^{2-\alpha} - u_+^{2-\alpha}\Big), & 1 < \alpha < 2. \end{cases}$$

B.10.3 Mapping from initial to terminal wealth

Given the solved (s, t), compute the cut–offs $W_0^{\text{part}}, W_0^{\text{info}}$ and classify households by W_0:

$$\text{state}(W_0) = \begin{cases} \text{Non-participant} & \text{if } W_0 < W_0^{\text{part}}, \\ \text{Uninformed participant} & \text{if } W_0^{\text{part}} \le W_0 < W_0^{\text{info}}, \\ \text{Informed participant} & \text{if } W_0 \ge W_0^{\text{info}}. \end{cases}$$

For the inequality exercise we model terminal wealth as a piecewise linear scaling of W_0:

$$W_2(W_0) = \begin{cases} r\, W_0, & \text{Non-participant}, \\ (1+\delta_U)\, W_0, & \text{Uninformed}, \\ (1+\delta_I)\, W_0, & \text{Informed}, \end{cases}$$

where δ_U and δ_I are calibrated to match aggregate trading gains implied by the model,[20] or set as scenario constants (e.g. $\delta_U = 5\%$, $\delta_I = 10\%$).

B.11 Gini coefficient

For any distribution of terminal wealth W_2 with mean μ_2,

$$\text{Gini} = \frac{1}{2\mu_2}\, \mathbb{E}\big[\, \big|W_2 - W_2'\big| \,\big],$$

where W_2, W_2' are i.i.d. draws from the same distribution. Operationally (simulation with N draws), sort $\{W_2^{(i)}\}_{i=1}^N$ ascending to $\{W_2^{(i:N)}\}$ and compute

$$\widehat{\text{Gini}} = \frac{1}{\bar{W}_2\, N\,(N-1)} \sum_{i=1}^{N} \big(2i - N - 1\big)\, W_2^{(i:N)}, \qquad \bar{W}_2 = \frac{1}{N}\sum_{i=1}^{N} W_2^{(i)}.$$

B.12 Simulation design for inequality experiments

We study two policy shocks—*entry-fee cut* ($F \downarrow$) and *information-cost cut* ($c_1 \downarrow$)—under two initial wealth distributions:

1. **Log-normal**: $\log W_0 \sim \mathcal{N}(\mu, \sigma^2)$,

2. **Pareto**: $P(W_0 > w) = \big(w_{\min}/w\big)^{\alpha}$, $\alpha \in (1, 2]$.

[20] From Section 2, the expected trading gain for an uninformed (informed) participant is proportional to the expected squared Sharpe ratio; choosing δ_U and δ_I to match those moments keeps the mapping consistent with the theory while keeping the inequality accounting transparent.

Parameter baseline. $\rho = 0.90$, $\sigma_z^2 = 1$, $\sigma_x^2 = 1$, $h_0 = 0.3$, $r = 0.07$, $c_0 = 0.2$, $F = 2$, $c_1 = 2$. The targets for baseline shares of participating and informed investors are $= 0.60$, and $= 0.10$ respectively. Shocks: $F \to 0.6$ (70% entry-fee cut) and $c_1 \to 0.6$ (70% info-cost cut).

Algorithm.

1. For baseline and each shock, solve (B.34)-(B.35) for (s, t) by Newton's method. Solve also $D(s, t)$, W_0^{part}, W_0^{info} via (10).

2. Compute $W_0^{\text{part}}, W_0^{\text{info}}$ and assign each $W_0^{(i)}$ to a state.

$$\text{state} = \begin{cases} \text{non-participant}, & W_0 < W_0^{\text{part}}, \\ \text{uninformed part.}, & W_0^{\text{part}} \le W_0 < W_0^{\text{info}}, \\ \text{informed part.}, & W_0 \ge W_0^{\text{info}}. \end{cases}$$

3. Map to terminal wealth $W_2^{(i)}$ via the piecewise rule above.

4. Compute $\widehat{\text{Gini}}$.

Note that:
(i) The Jacobian $J(\theta)$ is well-conditioned in all calibrations I tested (diagonal entries exceed 1 in typical cases), so the Newton step converges in a handful of iterations. (ii) For Pareto with $\alpha \le 2$ use the capped M_2 in Section 3.3 to avoid a divergent second moment—this is also realistic given finite top wealth in the data. (iii) Because $\sigma_{sj}^{-2}(W_{\text{info}}) = 0$, the lower-limit terms drop from Ψ_{2s}, Ψ_{2t}, which stabilises the solver numerically.

Expected outcomes.

- **Cut in** F: $W_0^{\text{part}} \downarrow$, $W_0^{\text{info}} \uparrow$ $\Rightarrow$ more passive entrants, less concentration of information $\Rightarrow$ Gini $\downarrow$ (especially for Log-normal).

- **Cut in** c_1: $W_0^{\text{info}} \downarrow$, $W_0^{\text{part}} \uparrow$ $\Rightarrow$ rich become informed, some middle-class exit $\Rightarrow$ Gini $\uparrow$, with the effect strongest under the Pareto tail.

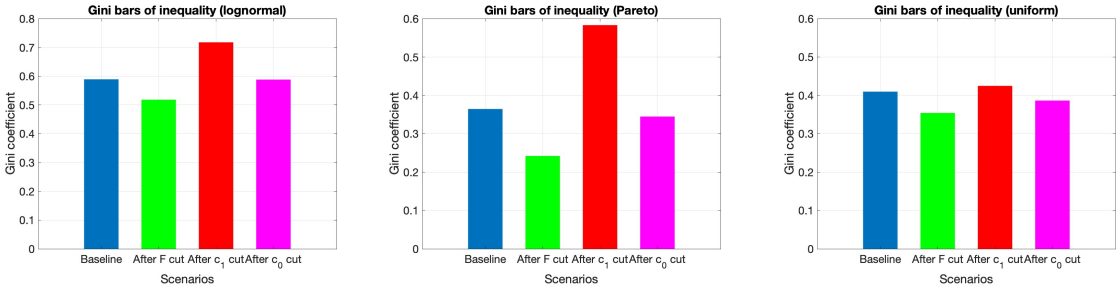
B.13 Simulation results

B.13.1 Plots of results under Lognormal, Pareto, and Uniform initial distributions

Gini inequality Figure 12 displays the Gini inequality index across all initial wealth distributions: Lognormal, Pareto, and Uniform.

Figure 12: **Gini across initial wealth distributions.** Left: Under the Lognormal initial wealth distribution. Middle: Under the Pareto initial wealth distribution. Right: Under the Uniform initial wealth distribution. The ranking of Gini inequality measures (Y-axis) stays the same across the scenarios (X-axis) independent of the initial wealth distributions (Panels).

Left. Log-normal Middle. Pareto Right. Uniform



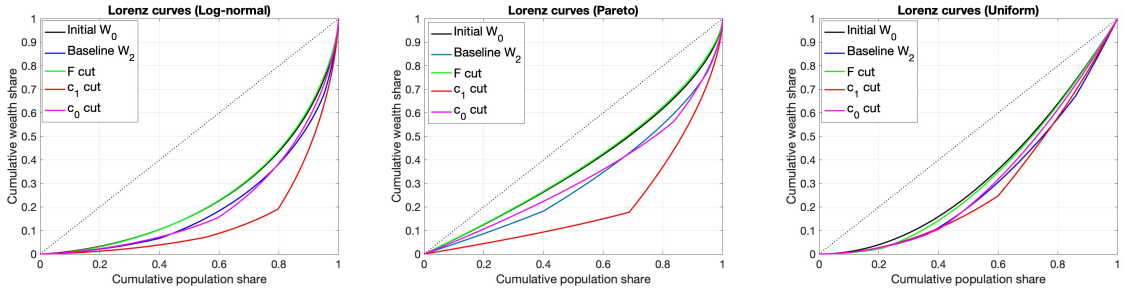
Notes. Initial wealth W_0 and primitives are held fixed across scenarios; in each case we solve the fixed point for (s, t). The axes show the Gini level on a common scale across panels. Patterns are consistent with the comparative statics independently of the distribution used: the F cut compresses inequality, the c_1 cut raises it, and the c_0 cut leaves it broadly unchanged. Variable definitions are in Appendix A.

Lorenz curves Figure 13 displays the Lorenz curves across all initial wealth distributions: Lognormal, Pareto, and Uniform.

PDFs of terminal wealth Figure 14 displays the terminal wealth distributions across all initial wealth distributions: Lognormal, Pareto, and Uniform. We superimpose the initial wealth distribution in gray on each graph.

Figure 13: **Lorenz curves across initial wealth distributions.** Left: Under the Lognormal initial wealth distribution. Middle: Under the Pareto initial wealth distribution. Right: Under the Uniform initial wealth distribution. The ranking of Lorenz inequality measures stays the same across the scenarios (in colors) independent of the initial wealth distributions (Panels).

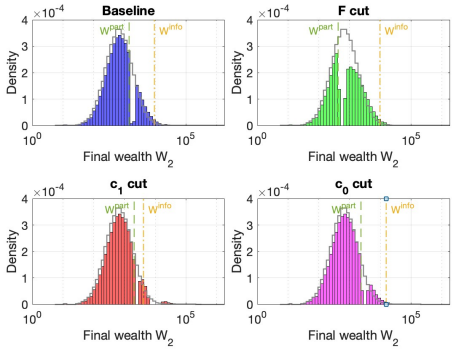
Left. Log-normal Middle. Pareto Right. Uniform



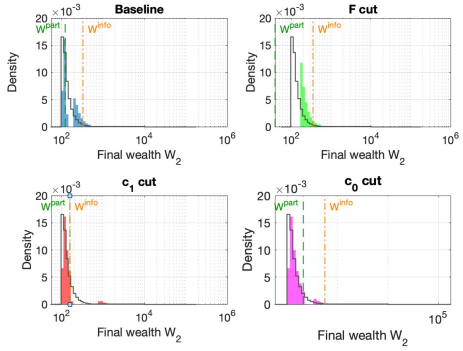
Notes. Initial wealth W_0 and primitives are held fixed across scenarios; in each case we solve the fixed point for (s, t). Patterns are consistent with the comparative statics independently of the distribution used: the F cut compresses inequality, the c_1 cut raises it, and the c_0 cut leaves it broadly unchanged. Variable definitions are in Appendix A.

Figure 14: **PDFs of terminal wealth across initial wealth distributions.** Top: Under the Lognormal initial wealth distribution. Middle: Under the Pareto initial wealth distribution. Bottom: Under the Uniform initial wealth distribution.

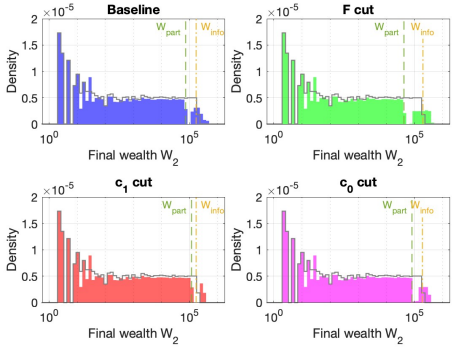
Left. Log-normal



Middle. Pareto



Right. Uniform



Notes. Initial wealth W_0 and primitives are held fixed across scenarios; in each case we solve the fixed point for (s,t). Patterns are consistent with the comparative statics independently of the distribution used: the F cut compresses inequality, the c_1 cut raises it, and the c_0 cut leaves it broadly unchanged. Variable definitions are in Appendix A.

B.14 Robustness

B.14.1 Alternative shares of participation and information.

Table 8: Gini Coefficient (Lognormal): Robustness for different levels of shares of participating and informed investors.

Share participation	90%	*60%*	30%	60%	60%	60%	60%	60%	60%	80%	90%	90%	99%
Share informed	5%	*5%*	5%	20%	10%	5%	2.50%	1%	0.10%	0.1%	0.1%	0.01%	0.01%
Baseline	0.559	*0.589*	0.622	0.604	0.599	0.589	0.580	0.571	0.562	0.539	0.530	0.528	0.523
F cut	0.522	*0.518*	0.530	0.513	0.515	0.518	0.521	0.523	0.525	0.523	0.522	0.522	0.522
c_1 cut	0.692	*0.718*	0.731	0.619	0.680	0.718	0.728	0.709	0.625	0.603	0.587	0.551	0.534
c_0 cut	0.537	*0.588*	0.602	0.594	0.590	0.588	0.587	0.586	0.584	0.553	0.536	0.535	0.523

Notes. The first two rows show the combinations of targeted shares for participating and informed investors. The last four rows show the Gini coefficient for a given scenario, keeping all else constant. Each scenario applies a one-at-a-time 70% cut to the named cost. The scenario in the main text is in Italics. This table is for the Lognormal prior wealth distribution. Variable definitions are in Appendix A.

Table 9: Gini Coefficient (Pareto): Robustness for different levels of shares of participating and informed investors.

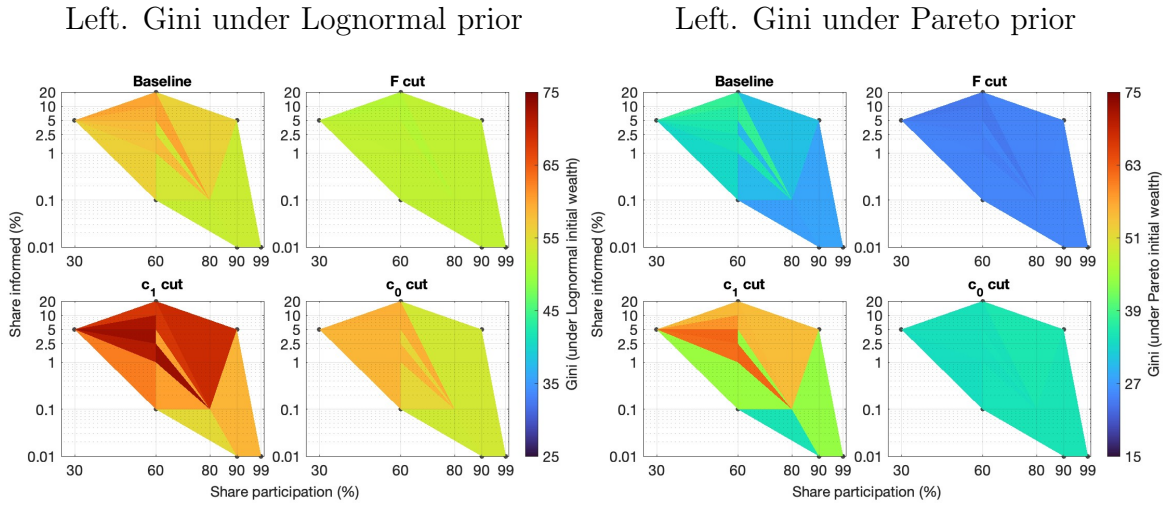
Share participation	90%	*60%*	30%	60%	60%	60%	60%	60%	60%	80%	90%	90%	99%
Share informed	5%	*5%*	5%	20%	10%	5%	2.50%	1%	0.10%	0.1%	0.1%	0.01%	0.01%
Baseline	0.308	*0.365*	0.388	0.390	0.377	*0.365*	0.355	0.345	0.335	0.300	0.279	0.276	0.254
F cut	0.250	*0.242*	0.231	0.233	0.238	*0.242*	0.245	0.248	0.250	0.251	0.251	0.251	0.251
c_1 cut	0.544	*0.583*	0.552	0.514	0.549	*0.583*	0.619	0.612	0.445	0.443	0.440	0.351	0.348
c_0 cut	0.354	*0.345*	0.324	0.349	0.346	*0.345*	0.344	0.343	0.342	0.349	0.351	0.351	0.353

Notes. The first two rows show the combinations of targeted shares for participating and informed investors. The last four rows show the Gini coefficient for a given scenario, keeping all else constant. Each scenario applies a one-at-a-time 70% cut to the named cost. The scenario in the main text is in Italics. This table is for the Pareto prior wealth distribution. Variable definitions are in Appendix A.

Tables 8 (for lognormal) and 9 (for Pareto) report the Gini coefficient of final wealth W_2 on a grid of participating rates $(30, 40, 80, 90, 99)$ and informed shares $(0.01, 0.1, 1, 2.5, 5, 10, 20)$ percent. Each cell corresponds to a recalibrated fixed point with $\mathbb{E}w_2$ held constant; infeasible pairs with informed share above participation are omitted. Figure 15 visualizes these cells as heatmaps using common color limits $[0.25, 0.75]$ for comparability. The Gini falls along the participation axis and rises along the informed axis for both priors; gradients are uniformly steeper under the Pareto prior, consistent with the role of tail thickness.

Figure 15 graphically plots the Gini coefficients under the Lognormal and the Pareto wealth distributions under varying shares of participating and informed investors. Comparing the left and right panels in Figure 15 highlights the role of tail

Figure 15: **Gini across initial wealth distributions.** Left: Under the Lognormal initial wealth distribution. Right: Under the Pareto initial wealth distribution. Each panel plots the Gini coefficient of final wealth W_2 as a function of the stock-market participation share (x-axis, %) and the informed-investor share (y-axis, %, log scale) in the Baseline and after one-at-a-time 70% cuts to F (participation cost), c_1 (marginal information cost), and c_0 (quadratic information cost). The mechanism is not an artifact of the lognormal prior. Repeating the exercise with a Pareto initial wealth distribution preserves the ranking: lowering the entry cost F compresses inequality, lowering the marginal information cost c_1 widens it, and lowering the quadratic cost c_0 does little.



Notes. Initial wealth W_0 and primitives are held fixed across scenarios; in each case we solve the fixed point for (s, t). Axes report model-implied population shares (percent). The colorbar shows the Gini level on a common scale across panels; filled squares mark the model-evaluated *(participation, informed)* combinations and the shading interpolates between them. Patterns are consistent with the comparative statics independently of the distribution used: the F cut compresses inequality, the c_1 cut raises it, and the c_0 cut leaves it broadly unchanged. Variable definitions are in Appendix A.

risk. Under lognormal W_0 , the iso-Gini contours are relatively flat in the informed dimension: raising the informed share increases the Gini, but moderately. Under the Pareto prior, the same move produces a much larger increase, and the downward gradient along participation is also stronger. Intuitively, heavy right tails make information more valuable at the top and make broad-based entry more compressing at the bottom. Fixing the color scale across panels makes these differences transparent.

B.14.2 Alternative magnitudes of technology improvement shocks

Table 10 presents the Gini coefficient across the two prior wealth distributions (i.e., Lognormal and Pareto) for different magnitudes of cuts to the named cost, keeping

Table 10: Gini Coefficient: Robustness for different levels of technology improvements.

	Lognormal					*Pareto*				
Cut magnitude →	20%	*30%*	50%	70%	80%	*20%*	30%	50%	70%	80%
Baseline	0.589	*0.589*	0.589	0.589	0.589	*0.365*	0.365	0.365	0.365	0.365
F cut	0.570	*0.560*	0.539	0.518	0.509	*0.279*	0.259	0.251	0.242	0.236
c_1 cut	0.613	*0.629*	0.671	0.718	0.720	*0.407*	0.438	0.527	0.517	0.611
c_0 cut	0.586	*0.584*	0.584	0.588	0.594	*0.371*	0.371	0.364	0.345	0.328

Notes. The tables show the Gini coefficient across the two prior wealth distributions for different magnitudes of cuts to the named cost, keeping all else constant. Each scenario applies a one-at-a-time cut of magnitude presented in Row 2 to the named cost in Column 1. Variable definitions are in Appendix A.

all else constant. Each scenario applies a one-at-a-time cut of magnitude presented in Row 2 to the named cost in Column 1. The Gini preserves the ranking: lowering the entry cost F compresses inequality, lowering the marginal information cost c_1 widens it, and lowering the quadratic cost c_0 does little across the different magnitudes of technology improvement shocks.

B.14.3 Alternative initial wealth distributions

We repeat the baseline four-scenario exercise—baseline calibration and one-at-a-time cuts in F, c_1, and c_0—under alternative priors for W_0. For each prior, we recalibrate F, c_1 to match the same baseline participation (60%) and informed share (5%), and we re-scale to hold $\mathbb{E}W_2$ fixed. We then solve the two fixed-point conditions in (s, t), compute the induced wealth cutoffs W_0^{part} and W_0^{info} (whose ratio equals c_1/F), map W_0 into W_2, and report Gini and Lorenz curves. Table 11 lists (s, t), thresholds, shares, and Ginis for Lognormal, Pareto, and Uniform prior wealth distributions. Relative to the Lognormal, the Pareto prior (i) magnifies the Gini increase from a c_1 cut—because more mass is poised near the upper tail—and (ii) strengthens the Gini decline from an F cut—because entry draws in a thicker base of low-wealth households. c_1 cut moves both thresholds to the right proportionally, leaving inequality nearly unchanged in both priors. The smallest movements in the Gini (although the ranking is the same across the scenarios) are for the Uniform distribution. This is because the Uniform prior has no tails and a hard cap on wealth, so it cannot remove very rich or very poor households. This mechanically damps how much the Gini can move, although the direction of change (decline from an F cut, increase from a c_1 cut, minimal change from a c_0 cut) stays the same.

Table 11: **Aggregate (risk-weighted) precision and risk tolerance, wealth thresholds, shares of participating and informed investors, Gini coefficients across all prior wealth distributions.**

	Lognormal distribution							
Scenario ↓	s	t	D(e-12)	Wp	Wi	share_part	share_info	Gini
baseline	675.19	2964.3	2.673	1403.4	9366.1	60.00%	5.00%	0.59
F_cut	660.1	3288.2	2.641	426.09	9478.9	92.60%	4.90%	0.52
c1_cut	787.13	2664.5	1.8	2084.6	4173.8	44.30%	20.10%	0.72
c0_cut	820.92	2568.7	1.63	2304.9	15383	40.40%	1.60%	0.59
	Pareto distribution							
Scenario ↓	s	t	D(e-10)	Wp	Wi	share_part	share_info	Gini
baseline	303.44	136.3	D=5.37e-10	122.67	331.44	60.00%	5.00%	0.37
F_cut	279.9	185.19	D=5.088e-10	38.84	349.78	100%	4.40%	0.24
c1_cut	375.42	92.363	D=4.143e-10	159	128.88	31.40%	53.00%	0.52
c0_cut	464.92	62.65	D=3.198e-10	205.96	556.5	16.40%	1.40%	0.35
	Uniform distribution							
Scenario ↓	s	t	D	Wp	Wi	share_part	share_info	Gini
baseline	0.3028	91420	2.72	72442	172079	63.80%	14.00%	0.41
F_cut	0.3251	106188	2.36	41685	198036	79.20%	1.00%	0.35
c1_cut	0.383	70924	1.64	119952	172079	40.00%	0.30%	0.43
c0_cut	0.3187	91420	2.46	80067	190193	60.00%	4.90%	0.39

Notes. The tables show the Gini coefficient across the two prior wealth distributions for different magnitudes of cuts to the named cost, keeping all else constant. Each scenario applies a one-at-a-time cut of magnitude presented in Row 2 to the named cost in Column 1. Variable definitions are in Appendix A.

C Proof of $D(s,t) = S_U^2(s,t)$

Consider an *uninformed* participant ($\pi = 0$). Write $\tilde{z} \equiv z - \mu_z$ and $\tilde{x} \equiv x - \mu_x$ so that the linear price is

$$rp \;=\; a \;+\; b\,\tilde{z} \;-\; c\,\tilde{x}, \qquad \mathbb{V}(rp) = b^2\sigma_z^2 + c^2\sigma_x^2. \tag{C.1}$$

Posterior moments conditional on p. Joint normality implies

$$\mathbb{E}[\tilde{z} \mid p] = \kappa\,(rp - \mathbb{E}[rp]), \quad \kappa = \frac{Cov(\tilde{z}, rp)}{\mathbb{V}(rp)} = \frac{b\sigma_z^2}{b^2\sigma_z^2 + c^2\sigma_x^2},$$

and

$$\mathbb{V}(\tilde{z} \mid p) = \frac{1}{h_0 + h_p}, \qquad h_p = \frac{b^2}{c^2\sigma_x^2} = \frac{s^2}{\sigma_x^2}, \quad \bar{h} \;\equiv\; h_0 + h_p. \tag{C.2}$$

Excess payoff and Sharpe ratio. The risky excess payoff is $z - rp = \tilde{z} - (rp - \mathbb{E}[rp]) - \mu_x/t$, so *conditional* excess return is

$$\mathbb{E}[z - rp \mid p] = (\kappa - 1)(rp - \mathbb{E}[rp]) - \frac{\mu_x}{t}.$$

Define the conditional Sharpe ratio

$$\eta \;=\; \frac{\mathbb{E}[z - rp \mid p]}{\sqrt{\mathbb{V}(\tilde{z} \mid p)}} \;=\; \big(\kappa - 1\big)\,\sqrt{\bar{h}}\,(rp - \mathbb{E}[rp]) \;-\; \frac{\mu_x}{t}\,\sqrt{\bar{h}}. \tag{C.3}$$

Expected squared Sharpe ratio. Taking expectations over p and using $\mathbb{E}[(rp - \mathbb{E}[rp])^2] = \mathbb{V}(rp)$:

$$S_U^2(s,t) \;\equiv\; \mathbb{E}[\eta^2] = (\kappa - 1)^2\,\bar{h}\,\mathbb{V}(rp) + \frac{\mu_x^2}{t^2}\bar{h}.$$

$Insert\kappa$ and $\mathbb{V}(rp)$ from (C.1)–(C.2)

$$=\Big[\frac{\big(b\sigma_z^2 - \mathbb{V}(rp)\big)^2}{\mathbb{V}(rp)} + \frac{\mu_x^2}{t^2}\Big].$$

Using the equilibrium identity $b = s\,c$ and $\mu_x = -t\,(\mu_z - \mathbb{E}[rp])$, straight algebra simplifies the bracket to $\sigma_x^2 + 2st + t^2 h_0$. Also $\bar{h} = h_0 + s^2/\sigma_x^2$. Hence

$$\boxed{S_U^2(s,t) = \frac{\sigma_x^2 + 2st + t^2 h_0}{t^2\big(h_0 + s^2/\sigma_x^2\big)^2}}. \tag{C.1}$$

Equality with $D(s,t)$. Section 3 defined the marginal gain from an infinitesimal precision,

$$D(s,t) \;=\; \frac{\sigma_x^2 + 2st + t^2 h_0}{t^2\big(h_0 + s^2/\sigma_x^2\big)^2},$$

which is *identical* to (C.1). Therefore

$$\boxed{D(s,t) \;=\; S_U^2(s,t) \quad \text{for all } s > 0,\ t > 0.} \tag{C.2}$$

This proves that the marginal benefit of producing a bit of private information equals the expected squared Sharpe ratio available to an uninformed participant—exactly the analogue of Peress's "A equals $w(0)$" result, but under our quadratic-plus-linear cost specification.

D Delegation: Details and Proofs

D.1 Manager's problem and cutoffs

For a given (s,t) and pool size $T_{\mathcal{M}}$, the manager chooses precision h to maximize $\frac{T_{\mathcal{M}}}{2}D(s,t)\,h - c_1 h - \frac{c_0}{2}h^2$, yielding $h_m = (T_{\mathcal{M}}D - 2c_1)/(2c_0) \vee 0$ and the pool-level information threshold $T_{\mathcal{M}}^{\text{info}} = 2c_1/D(s,t)$. Investor j delegates when $\frac{T_j}{2}D\,h_m \geq f + F_m$, that is, $W_{0j} \geq \rho \cdot 2(f+F_m)/(D\,h_m)$.

D.2 Fees and industry zero-profit

With per-capita fees, competitive managers break even in the cross section: $f = C(h_m)/M = c_1 h_m/M + \frac{c_0}{2}(h_m)^2/M$. Other fee schedules produce the same local cutoffs when linearized around h_m.

D.3 Delegation: Fees, Matching, and Pooling Assumptions

We allow investors to pool and delegate information acquisition to a competitive manager. Managers are technologically identical, face the same convex information technology $C(h) = c_1 h + \frac{c_0}{2}h^2$ with $c_0, c_1 > 0$, and observe an asset-level signal of precision $h \geq 0$ on behalf of their pool. Matching is frictionless: any investor may join any fund. Let $\mathcal{M}$ denote the endogenous set of delegators and $T_{\mathcal{M}} \equiv \int_{j\in\mathcal{M}}(W_{0j}/\rho)\,dV$ its *aggregate risk tolerance*. We assume managers are *price takers*: competition drives fees to (at least) cover information costs.[21]

Fee setting. In the baseline, competition implies zero rents for managers: the per-capita fee equals average cost,

$$f \;=\; \frac{C(h_m)}{M} \qquad \text{with} \qquad h_m \in \arg\max_{h\geq 0}\Big\{\frac{T_{\mathcal{M}}}{2}D(s,t)\,h - C(h)\Big\},$$

where $M \equiv \text{mass}(\mathcal{M})$ and $D(s,t)$ is the marginal gain per unit precision defined in the baseline. An optional *bargaining* variant lets the manager capture a share $\eta \in [0,1]$ of joint surplus $JS = \frac{c_0}{2}(h_m)^2$, in which case $f = \frac{C(h_m)}{M} + \eta\,\frac{JS}{M}$. The baseline corresponds to $\eta = 0$.

[21] This is the delegated analogue of "alpha equals costs" in efficiently inefficient markets; see Gârleanu and Pedersen (2018) for microfoundations and related delegated-management environments in Cuoco and Kaniel (2011); He and Xiong (2013).

Pooling and precision faced by delegators. Delegators share the *same* signal with precision h_m. Individual precision *does not depend* on own wealth; instead it scales with pooled risk tolerance:

$$h_m(T_{\mathcal{M}}; s, t) = \left(\frac{T_{\mathcal{M}} D(s,t) - 2c_1}{2c_0}\right)_+.$$

In price aggregation, the delegated block contributes $T_{\mathcal{M}} h_m$ to s, so the price rule remains $rp = a + bz - cx$ with $b = sc$ and precision $h_p = s^2/\sigma_x^2$ as in the baseline.

D.4 Self-learning

Proposition 6 (Self dominates delegation at the top)**.** *Assume $c_1 > F$ and $T_{\mathcal{M}} D > 2c_1$ so that $W_{\text{info}} > W_{\text{part}}$ and $h_m > 0$. If $\Delta_{\text{del}}(W_{\text{info}}) > 0$ (so the delegation window opens), then there exists a unique $W_{\text{switch}} > W_{\text{info}}$ such that*

$$\Delta_{\text{self}}(W) \;<\; \Delta_{\text{del}}(W) \quad \textit{for } W \in [W_{\text{info}}, W_{\text{switch}}), \qquad \Delta_{\text{self}}(W) \;>\; \Delta_{\text{del}}(W) \quad \textit{for } W > W_{\text{switch}}.$$

Equivalently, writing $X \equiv (W/\rho)D$, $W_{\text{switch}} = (\rho/D)\, X_{\text{switch}}$ where

$$X_{\text{switch}} \;=\; 2\big(c_1 + c_{0s} h_m\big) \;+\; 2\sqrt{c_{0s}\Big(2c_1 h_m + c_{0s}(h_m)^2 - 2(f + F_m)\Big)}.$$

Proof. Note that $\Delta_U \text{self}(W) = \big(X - 2c_1\big)_+^2/(8c_{0s})$ is convex and increasing for $X \geq 2c_1$, while $\Delta_{\text{del}}(W) = \frac{1}{2} X h_m - f - F_m$ is affine in X.

If the window opens, the two curves intersect once for $X \geq 2c_1$, yielding the stated root.

For $X > X_{\text{switch}}$ the self curve grows quadratically while the delegation curve grows linearly, hence self dominates for sufficiently high W.

This concludes the proof. □

D.5 Market clearing with delegation and fixed-point algorithm

Let $\mathcal{U}, \mathcal{I}, \mathcal{M}$ be the endogenous sets of uninformed participants, self-informed participants, and delegators determined by the cutoffs $W_0^{\text{part}}, W_0^{\text{del}}, W_0^{\text{info}}, W_0^{\text{switch}}$. Aggregate

objects are

$$t = \int_{\mathcal{U}\cup\mathcal{I}\cup\mathcal{M}} T_j \, dV, \qquad s = \int_{\mathcal{I}} T_j h_j \, dV + T_{\mathcal{M}} h_m,$$

and the price rule is $rp = a + bz - cx$ with $b = sc$ and $h_p = s^2/\sigma_x^2$ as in the baseline. The fixed point in (s, t) extends the map in Appendix B.5 by composing it with the inner map $(h_m, M) \mapsto (T_{\mathcal{M}}, f)$ defined by the fee and cutoff conditions. Continuity of the composition together with interior signs $D_s < 0$ and $D_t < 0$ imply a nonsingular Jacobian at the solution and local convergence of Newton iterations started in a neighborhood.

Economic intuition. The manager chooses a *single* precision $h_m > 0$ to suit the pool's average tolerance $T_{\mathcal{M}}$, whereas investor j's idiosyncratic optimum is $h_j(W) = \big((W/\rho)D - 2c_1\big)/(2c_{0s})$, which scales linearly with wealth. Because the net self gain Δ_{self} grows *quadratically* with W once feasible, while the net delegation gain Δ_{del} grows only *linearly*, the very rich eventually prefer to self–learn even if $c_{0m} \le c_{0s}$. A manager offering a full menu of (h, fee) by wealth would reduce this wedge, but that institutional feature is outside our baseline and is briefly discussed in the main text.

D.6 Proofs for the Delegation Extension

Proof of Lemma 4 (Left boundary and collapse). Fix aggregate objects (s, t) and take $D \equiv D(s, t)$ as given when comparing individual choices. Let $W_{\text{part}} = \frac{2\rho F}{D}$ and $W_{\text{info}} = \frac{2\rho c_1}{D}$ denote the baseline participation and self–learning cutoffs, and assume $c_1 > F$, so $W_{\text{info}} > W_{\text{part}}$. Suppose also $T_{\mathcal{M}} D > 2c_1$, so the manager's optimal pooled precision in (13) is $h_m > 0$. Throughout, write

$$\Delta U^{\text{del}}(W) \;=\; \frac{W}{\rho}\,\frac{D}{2}\,h_m \;-\; f \;-\; F_m, \qquad \Delta U^{\text{self}}(W) \;=\; \frac{\big(\frac{W}{\rho}D - 2c_1\big)_+^2}{8c_0},$$

and define the delegation window as the set of wealth levels $W \ge W_{\text{part}}$ such that $\Delta U^{\text{del}}(W) \ge \max\{0, \Delta U^{\text{self}}(W)\}$.

Step 1 (Shape and single crossing). For $W \ge W_{\text{info}}$, $\Delta U^{\text{self}}(W) = \big(\frac{W}{\rho}D - 2c_1\big)^2/(8c_0)$ is strictly convex in W with $\Delta U^{\text{self}}(W_{\text{info}}) = 0$ and right derivative 0 at W_{info}. By contrast, $\Delta U^{\text{del}}(W)$ is affine in W with positive slope

$$\alpha \;\equiv\; \frac{D}{2\rho}\,h_m \;>\; 0.$$

Hence, for $W \geq W_{\text{info}}$ the difference

$$G(W) \;\equiv\; \Delta U^{\text{del}}(W) - \Delta U^{\text{self}}(W)$$

is continuous and strictly concave, satisfies $G(W_{\text{info}}) = \Delta U^{\text{del}}(W_{\text{info}})$, and obeys $G(W) \to -\infty$ as $W \to \infty$ (quadratic term dominates linear). Therefore, whenever $G(W_{\text{info}}) \geq 0$ there exists a unique $W_{\text{switch}} > W_{\text{info}}$ solving $G(W_{\text{switch}}) = 0$, with $G(W) > 0$ on $[W_{\text{info}}, W_{\text{switch}})$ and $G(W) < 0$ on $(W_{\text{switch}}, \infty)$.

Step 2 (Left boundary at W_{info} when the window is nonempty). For $W < W_{\text{info}}$, $\Delta U^{\text{self}}(W) = 0$ and $\Delta U^{\text{del}}(W)$ is strictly increasing in W. Consequently, if $G(W_{\text{info}}) > 0$ then the (unique) zero of $\Delta U^{\text{del}}(W)$ occurs weakly to the left of W_{info}:

$$W_0^{\text{del}} \;=\; \inf\{W : \Delta U^{\text{del}}(W) \geq 0\} \;\leq\; W_{\text{info}}.$$

Because $W_{\text{part}} < W_{\text{info}}$ by $c_1 > F$, the left boundary of the delegation window,

$$W_{\text{del}} \;=\; \inf\big\{W \geq W_{\text{part}} : \Delta U^{\text{del}}(W) \geq \max\{0, \Delta U^{\text{self}}(W)\}\big\},$$

equals W_{info}: below W_{info} the relevant benchmark is 0 (not ΔU^{self}), while at W_{info} we have $\Delta U^{\text{self}} = 0$ and $\Delta U^{\text{del}} > 0$, so delegation strictly dominates. By Step 1 there is then a unique $W_{\text{switch}} > W_{\text{info}}$ with delegation strictly preferred on $[W_{\text{info}}, W_{\text{switch}})$ and self–learning strictly preferred for $W > W_{\text{switch}}$.

Step 3 (Collapse). If $h_m \leq 0$, then $\Delta U^{\text{del}}(W) \leq -f - F_m < 0$ for all W, so the window is empty. If $h_m > 0$ but $G(W_{\text{info}}) = \Delta U^{\text{del}}(W_{\text{info}}) \leq 0$, concavity of G together with $G(W) \to -\infty$ as $W \to \infty$ implies that $G(W) < 0$ on $[W_{\text{info}}, \infty)$ (otherwise a nonempty window would require $G(W_{\text{info}}) \geq 0$ by Step 1 and the intermediate value property). Hence no $W \geq W_{\text{part}}$ satisfies $\Delta U^{\text{del}}(W) \geq \max\{0, \Delta U^{\text{self}}(W)$, and the delegation window is empty. By convention we then report $W_{\text{del}} = W_{\text{switch}}$.

Combining Steps 1–3 proves the lemma. □

Proof. Fix (s, t) and a pool $(T_{\mathcal{M}}, M)$.

(i) Manager's problem and thresholds. The manager chooses $h \geq 0$ to maximize

$$\Psi(h) \;=\; \frac{T_{\mathcal{M}}}{2} D(s,t)\, h \;-\; c_1 h \;-\; \frac{c_0}{2} h^2.$$

Ψ is strictly concave because $c_0 > 0$. The FOC and complementary slackness yield

$$h_m = \Big(\frac{T_{\mathcal{M}} D(s,t) - 2c_1}{2c_0}\Big)_+, \qquad T_{\mathcal{M}}^{\text{info}} = \frac{2c_1}{D(s,t)}.$$

The joint surplus at the optimum is $JS = \Psi(h_m) = \frac{c_0}{2}(h_m)^2$.

(ii) Fees. With competitive managers ($\eta = 0$), zero profits imply per-capita fees $f = \frac{C(h_m)}{M} = \frac{c_1 h_m + \frac{c_0}{2}(h_m)^2}{M}$. With Nash bargaining weight $\eta \in [0,1]$, the per-capita fee becomes $f = \frac{C(h_m)}{M} + \eta \frac{JS}{M}$.

(iii) Delegation cutoff and comparative statics (partial). Investor j's incremental mean–variance gain from delegation is $\Delta U_j^{\text{del}} = \frac{T_j}{2} D(s,t)\, h_m - (f + F_m)$ with $T_j = W_{0j}/\rho$. The cutoff W_0^{del} solves $\Delta U_j^{\text{del}} = 0$, hence $W_0^{\text{del}} = \rho \frac{2(f+F_m)}{D(s,t)h_m}$. Holding (s,t,M) fixed,

$$\frac{\partial W_0^{\text{del}}}{\partial F_m} = \frac{2\rho}{D(s,t)h_m} \;>\; 0.$$

Using $h_m = \frac{T_{\mathcal{M}} D - 2c_1}{2c_0}$, we have $\partial h_m/\partial c_1 = -1/c_0$ and $\partial h_m/\partial c_0 = -h_m/c_0$. Since $f = \frac{c_1 h_m + \frac{c_0}{2}(h_m)^2}{M}$, direct differentiation gives

$$\frac{\partial W_0^{\text{del}}}{\partial c_1} = \frac{2\rho}{c_0 D}\Big(\frac{1}{2}\, c_0\, \frac{1}{M} + \frac{F_m}{h_m}\Big) \;>\; 0, \qquad \frac{\partial W_0^{\text{del}}}{\partial c_0} = \frac{2\rho}{D}\Big(\frac{F_m}{c_0 h_m} + \frac{1}{2}\,\frac{1}{M}\Big) \;>\; 0.$$

Thus, ceteris paribus, higher c_1, c_0, or F_m raise the delegation cutoff.

(iv) Pooling superadditivity. For $T \geq T_{\mathcal{M}}^{\text{info}}$, $T h_m(T) = \frac{T(TD - 2c_1)}{2c_0}$ is convex in T. Hence, for any T_1, T_2 in the interior region,

$$T_1 h_m(T_1) + T_2 h_m(T_2) \;=\; \frac{D(T_1^2 + T_2^2) - 2c_1(T_1 + T_2)}{2c_0}$$

$$\leq\; \frac{D(T_1 + T_2)^2 - 2c_1(T_1 + T_2)}{2c_0} \;=\; (T_1 + T_2) h_m(T_1 + T_2),$$

and the inequality is strict whenever $T_1, T_2 > 0$. If one pool is below threshold while $T_1 + T_2$ exceeds it, pooling strictly improves s by turning zero into positive precision. Hence a representative fund spans investors without loss.

(v) Price rule and fixed point. The delegated block adds $T_{\mathcal{M}} h_m$ to s, so the price rule $rp = a + bz - cx$ with $b = sc$ and $h_p = s^2/\sigma_x^2$ is unchanged in form. Existence and local uniqueness of the outer fixed point in (s,t) follow from the baseline argument (continuity and $D_s < 0, D_t < 0$ on the interior region), since the inner map $(T_{\mathcal{M}}, M) \mapsto T_{\mathcal{M}} h_m$ is continuous. □

D.7 Efficiency vs welfare vs inequality: an accounting decomposition

Let $Y(W_0; \theta)$ denote the certainty–equivalent (CE) terminal wealth for an investor with initial wealth W_0 under parameter vector $\theta \in \{F, F_m, c_1, c_0\}$ and equilibrium

objects (s, t, h_m, f).[22] Define utilitarian social welfare as

$$\mathrm{SW}(\theta) \equiv \int Y(W_0; \theta)\, \mathrm{d}V(W_0).$$

For a small policy change $\mathrm{d}\theta$, a first–order decomposition (Leibniz rule) yields

$$\begin{aligned}
\mathrm{dSW} = & \underbrace{\Big[Y^{\mathrm{mkt}}(W_0^{\mathrm{part}}) - Y^{\mathrm{out}}\Big]\big(-v(W_0^{\mathrm{part}})\,\mathrm{d}W_0^{\mathrm{part}}\big)}_{\text{(i) Participation margin}} + \\
& + \underbrace{\Big[Y^{\mathrm{del}}(W_0^{\mathrm{del}}) - Y^{\mathrm{uninf}}(W_0^{\mathrm{del}})\Big]\big(v(W_0^{\mathrm{del}})\,(-\mathrm{d}W_0^{\mathrm{del}})\big)}_{\text{(ii) Uninformed}\to\text{Delegation}} - \\
& - \underbrace{\Big[Y^{\mathrm{self}}(W_0^{\mathrm{switch}}) - Y^{\mathrm{del}}(W_0^{\mathrm{switch}})\Big]\big(v(W_0^{\mathrm{switch}})\,(-\mathrm{d}W_0^{\mathrm{switch}})\big)}_{\text{(iii) Delegation}\to\text{Self}} \\
& + \underbrace{\int_{\text{participants}} \frac{\partial Y}{\partial(s,t)} \cdot (\mathrm{d}s, \mathrm{d}t)\, \mathrm{d}V(W_0)}_{\text{(iv) Price-informativeness / risk-sharing}} - \underbrace{\mathrm{d}\big[\text{resource costs } C(h_j) + M \cdot C(h_m)\big]}_{\text{(v) Info production costs}},
\end{aligned}$$

where Y^{out} is the outside option CE, and $Y^{\mathrm{mkt}}, Y^{\mathrm{uninf}}, Y^{\mathrm{del}}, Y^{\mathrm{self}}$ are CEs in each regime. Term (iv) captures the *efficiency* channel: more s improves price informativeness but lowers the marginal value of information $D(s,t)$ ($D_s < 0$), which tightens W_0^{part} via adverse selection;[23] terms (i)–(iii) trace the *distribution* channel; term (v) nets out real resource use.

Implications. *Participation cuts* ($F \downarrow$) unambiguously raise SW and compress inequality: (i) is positive and large, while (iv)–(v) are second order locally because s falls only modestly when more risk is shared.[24] *Information-cost cuts* ($c_1 \downarrow$) raise s and efficiency (iv), but they push W_0^{part} right and reallocate mass toward informed types, making (i) negative and (ii)–(iii) inequality-increasing; the SW sign is calibration-dependent and—importantly—can be negative when the entry deterrence dominates. *Quadratic cost cuts* ($c_0 \downarrow$) mostly wash out in SW because they raise precisions but also lower $D(s,t)$; their inequality impact is theoretically ambiguous and small in the calibration.

[22] Notation as in Appendix D.7.

[23] Price rule $r_p = a + bz - cx$ with $b = sc$ and $h_p = s^2/\sigma_x^2$; the outer fixed point has $D_s < 0$, $D_t < 0$.

[24] Consistent with the model's ranking and simulated Gini movements when lowering F.

D.8 Comparative statics with delegation (costs and search).

Let $W_0^{\text{part}} = \frac{2\rho F}{D(s,t)}$, $W_0^{\text{info}} = \frac{2\rho c_1}{D(s,t)}$, and let the *delegation cutoff* for investor j be

$$W_0^{\text{del}} \;=\; \rho\,\frac{2\,(f+F_m)}{D(s,t)\,h_m},$$

where $F_m \geq 0$ is the search cost to locate a manager. Holding (s,t) fixed:

1. $F_m \downarrow \Rightarrow\; W_0^{\text{del}} \downarrow$ (direct effect), expanding delegation.
2. $c_1 \downarrow \Rightarrow\; h_m \uparrow$ and $W_0^{\text{del}} \downarrow$; the pool-level information threshold $T_{\mathcal{M}}^{\text{info}} = 2c_1/D$ also falls.
3. $c_0 \downarrow \Rightarrow\; h_m \uparrow$ and $W_0^{\text{del}} \downarrow$.
4. $F \downarrow \Rightarrow\; W_0^{\text{part}} \downarrow$; participation rises, t increases and (through $D_t < 0$) $D(s,t)$ falls in general equilibrium, which tends to move W_0^{info} to the right. The net effect on s is model-dependent because $T_{\mathcal{M}}$ and M also expand.
5. $\eta \uparrow$ (bargaining power of managers) $\Rightarrow f$ rises and $W_0^{\text{del}} \uparrow$, shrinking delegation.

Two composition remarks: (i) For *fixed* $T_{\mathcal{M}}$, increasing M lowers per-capita fees $f = C(h_m)/M$ (since h_m depends only on $T_{\mathcal{M}}$), thus reducing W_0^{del}. (ii) For *fixed* M, increasing $T_{\mathcal{M}}$ raises h_m and—by convex costs—raises $(f+F_m)/h_m$, thus increasing W_0^{del}. In equilibrium, $T_{\mathcal{M}}$ and M co-move; the empirical composition determines the net movement of W_0^{del}.

Corollary 1 (Inequality with delegation and cost cuts)**.** *Consider an interior equilibrium with a nonempty delegation window ($h_m > 0$ and $\Delta U^{\text{del}}(W_{\text{info}}) > 0$). Then the Gini coefficient G of terminal wealth exhibits the following local comparative statics:*

1. *Lower participation cost F* reduces *inequality: G strongly falls. W_0^{part} shifts left, bringing marginal (low-wealth) households from the outside option into the market. This* compresses *the cross-sectional distribution (Gini down). The new entrants are mostly uninformed at the margin; the informed share among participants changes little on impact. Any feedback through $D(s,t)$ is second-order locally, so on net, inequality falls.*
2. *Lower linear information cost c_1* raises *inequality: G weakly rises. A cheaper linear information cost boosts both self- learning and delegation, raises s, lowers D, and hence pushes W_0^{part}* right*. Some households exit, which* widens *the distribution and increases inequality strongly. At the same time, W_0^{info} shifts* left *and h_m rises ($h_m = \frac{T_{\mathcal{M}}D - 2c_1}{2c_0}$). A larger share becomes informed (self or delegated), increasing payoff dispersion within participants. The net effect is that inequality rises.*
3. *Lower quadratic information cost c_0 effect is calibration-dependent, but small: G*

can weakly rise or fall. Lower c_0 increases optimal precisions (self and pooled), raises s, reduces D, and thereby shifts W_0^{part} right *(harder entry). This* widens *the distribution (Gini up). For any given W, optimal individual self-learning precision h_j and the pool's precision h_m rise, steepening the payoff gradient with wealth (self's gain is quadratic in W, delegation's is linear). The right tail thickens. However, the pool can shrink too. On the net,*

4. *Lower quadratic information cost c_0 effect is calibration-dependent, but small: G can weakly rise or fall. Lower c_0 increases optimal precisions (self and pooled), raises s, reduces D, and thereby shifts W_0^{part}* right *(harder entry). This* widens *the distribution (Gini up). For any given W, optimal individual self-learning precision h_j and the pool's precision h_m rise, steepening the payoff gradient with wealth (self's gain is quadratic in W, delegation's is linear). The right tail thickens. On the net, the sign is calibration-dependent because both the participation margin and the informed composition move.*
5. *Lower delegation search cost F_m* raises *inequality: G weakly rises. Lower F_m expands delegation, raises the pool's precision, h_m, increases s, and thus lowers $D(s,t)$. Because $W_0^{part} = \frac{2\rho F}{D(s,t)}$, the fall in D shifts W_0^{part}* right *(adverse-selection channel), making entry harder. This* widens *the overall distribution and inequality increases. Moreover, as the pool's precision h_m increases, (and f falls with larger M), W_0^{del} moves* left*; more marginal participants switch from uninformed to delegated. Richer investors may switch from delegation to self (since ΔU^{self} grows quadratically in W). Both shifts fatten the right tail among participants. Overall, inequality G weakly rises.*

Remarks: (i) If the self–learning region is quantitatively large at baseline, the progressive self-to-delegated substitution triggered by $F_m \downarrow$ can partially offset the regressive uninformed→delegated expansion. In our calibration the latter dominates, so the net effect is inequality-increasing. (ii) However, if the delegation window collapses ($\Delta U^{\text{del}} \leq 0$), the F_m result is vacuous and the c_1, c_0 results revert to the baseline Corollary 1 (lower information costs increase G).

Under the same mean-preserving normalization as in the baseline, inequality reacts through two margins. (i) *Between-group* compression: if participation expands W_0^{part} declines), mass moves from the outside option into the market and the Gini tends to *decline*. (ii) *Within-participants* dispersion: if the informed share among participants rises (via self-learning or delegation), the right tail thickens and the Gini tends to *increase*.

In our extension, a cut in F_m, c_1, or c_0 all raise h_m and (weakly) the informed share via delegation; as there is a decrease in participation due to adverse selection, the total effect on the Gini is mostly upward. The quantitative impact is calibration-dependent only for the comparative static exercise with respect to a decrease in c_0 because the composition of delegated and self-informed investors also changes.

Proof of Corollary 1. Throughout fix an interior equilibrium with $h_m > 0$ and a nonempty delegation window. Let $Y(W;\theta)$ denote the certainty–equivalent (CE) terminal wealth of an investor with initial wealth W under parameter vector $\theta \in \{F, F_m, c_1, c_0\}$ and induced equilibrium objects (s, t, h_m, f).[25] We write $G(\theta)$ for the Gini coefficient of the (renormalized) distribution of $Y(\cdot;\theta)$.

We know that $W_{\text{part}} = \frac{2\rho F}{D(s,t)}$, $W_{\text{info}} = \frac{2\rho c_1}{D(s,t)}$, and $W_{\text{del}} = \rho \frac{2(f+F_m)}{D(s,t)\, h_m}$ are the participation, self–learning, and delegation cutoffs; W_{switch} is the unique solution to $\Delta U^{\text{del}}(W) = \Delta U^{\text{self}}(W)$ (Lemma 4). For $W \geq W_{\text{part}}$,

$$\Delta U^{\text{del}}(W) = \frac{W}{\rho}\frac{D}{2}h_m - (f + F_m), \qquad \Delta U^{\text{self}}(W) = \frac{(\frac{W}{\rho}D - 2c_1)_+^2}{8c_0}.$$

Hence, for participants, $Y(W;\theta) = W + \max\{0, \Delta U^{\text{del}}(W), \Delta U^{\text{self}}(W)\}$; for $W < W_{\text{part}}$, $Y(W;\theta) = W$ (outside option).

Two properties are used repeatedly:

P1 (Monotonicity and convexity in W). For fixed (s, t, h_m, f), $\Delta U^{\text{del}}(W)$ is affine and strictly increasing in W; $\Delta U^{\text{self}}(W)$ is strictly convex and strictly increasing on $[W_{\text{info}}, \infty)$ and zero below W_{info}. Therefore $Y(W;\theta)$ is increasing in W, and on any region where the identity of the optimal action does not change, it is (weakly) convex in W; when self–learning is optimal it is strictly convex.

P2 (Cutoff comparative statics). W_{part} and W_{info} move in the *same* direction as D^{-1}:

$$\frac{\partial W_{\text{part}}}{\partial D} = -\frac{2\rho F}{D^2} < 0, \qquad \frac{\partial W_{\text{info}}}{\partial D} = -\frac{2\rho c_1}{D^2} < 0.$$

Since $s = s_s + T_{hm}$ and $D = D(s,t)$ is strictly decreasing in s (baseline), any change that raises s lowers D and hence shifts both W_{part} and W_{info} to the right. Moreover, W_{del} is strictly increasing in $f + F_m$ and strictly decreasing in Dh_m.

We compare $G(\theta)$ under four local perturbations by decomposing the impact into (A) a *between–group* component driven by a change in W_{part} (entry/exit of low–W

[25] As in the baseline, we evaluate inequality on CE terminal wealth to stay within a static cross–sectional comparison in a CARA–normal environment. The mean–preserving normalization subtracts a scalar $\kappa(\theta)$ so that $\int Y(W;\theta)\, dF(W)$ is constant across scenarios.

households) and (B) a *within–participants* component driven by changes in actions (uninformed $\leftrightarrow$ delegated $\leftrightarrow$ self) and in optimal precisions. The Gini respects the Dalton principle of transfers; we repeatedly use that any progressive (regressive) mean–preserving transfer weakly lowers (raises) G.

Case F$\downarrow$ (cheaper participation). A decrease in F shifts W_{part} to the left holding D fixed at impact; feedback via D is second order locally. *Between–group:* new (lowest–W) entrants move from $Y = W$ to $Y = W + \max\{0, \Delta U^{\text{del}}, \Delta U^{\text{self}}\} \geq W$, which is a progressive transfer at the bottom; after mean normalization, G weakly falls. *Within–participants:* the identity of optimal actions and h_m are unchanged to first order; G is unaffected to first order. Hence $dG/dF < 0$ locally.

Case $F_m \downarrow$ (cheaper delegation search). Lower F_m (and the induced rise in M) increases the manager's optimal precision and lowers fees: $h_m \uparrow$, $f \downarrow$. Since s includes T_{hm}, we have $s \uparrow \Rightarrow D \downarrow$; thus $W_{\text{part}} \uparrow$ and $W_{\text{info}} \uparrow$ by P2. *Between–group:* the right shift of W_{part} excludes a marginal block of low–W households. This is a regressive change (mass leaves the bottom), so—after mean normalization—G weakly rises. *Within–participants:* two submargins move:

(i) Uninformed $\rightarrow$ Delegated: W_{del} falls because $F_m \downarrow$ and $f \downarrow$ (denominator Dh_m moves less at impact than $f + F_m$). Newly switched investors obtain an incremental payoff $\Delta U^{\text{del}}(W)$ which is strictly increasing in W; holding the mean fixed, this is a regressive perturbation on the participants' Lorenz curve and raises within–group inequality.

(ii) Self $\rightarrow$ Delegated: since W_{info} shifts right with $D \downarrow$, the self region shrinks. Because $\Delta U^{\text{self}}(W) - \Delta U^{\text{del}}(W)$ is strictly increasing in W (quadratic vs. linear), moving mass from self to delegated is a *progressive* adjustment within the right tail and therefore lowers within–group inequality.

The within–participants effect is thus the (positive) contribution from (i) minus the (negative) contribution from (ii). Locally, (i) dominates whenever the switch at the lower margin is larger than at the upper margin, which holds under the mild condition that the elasticity of the pool size M with respect to F_m exceeds the (small) induced shift in W_{info} via D:[26]

$$\textbf{(A1)} \qquad f_W(W_{\text{del}}) \left|\partial W_{\text{del}}/\partial F_m\right| \;\geq\; f_W(W_{\text{info}}) \left|\partial W_{\text{info}}/\partial F_m\right|,$$

[26]This condition is satisfied in calibrations where delegation is the extensive margin—M responds directly to F_m—and D is a smooth function of s so that the W_{info} shift is second order.

where f_W is the density of W. Under (A1), the within–participants effect is regressive and G weakly rises. Even without (A1), the between–group term is regressive, so G does not fall. Hence the stated sign in the corollary.

Case $c_1 \downarrow$ (cheaper linear information). Directly, $W_{\text{info}} = \frac{2\rho c_1}{D}$ shifts left; $h_m = \frac{T_D - 2c_1}{2c_0}$ rises; thus $s \uparrow$ and $D \downarrow$, which pushes W_{part} right (P2). *Between–group:* the rise in W_{part} is regressive $\Rightarrow$ G rises. *Within–participants:* both margins expand toward more information: (i) the self region grows (left shift of W_{info}), and (ii) delegation expands via $h_m \uparrow$ and $f \downarrow$. By P1, these shifts increase the convexity of $Y(W;\theta)$ on $[W_{\text{info}}, \infty)$ (self) and raise its slope on $[W_{\text{del}}, W_{\text{switch}})$ (delegation). With mean normalization, this is a regressive (Lorenz-worsening) change. Therefore G weakly rises.

Case $c_0 \downarrow$ (flatter marginal costs). Lower c_0 increases optimal precisions both individually and in the pool; hence $s \uparrow$ and $D \downarrow$ so that W_{part} shifts right (regressive). Within participants, $Y(W;\theta)$ becomes more convex in W (self's quadratic gain becomes larger, delegation's slope increases because $h_m \uparrow$), which is regressive after mean normalization. Therefore G weakly rises.

Quantitatively, the induced changes in type shares are

$$\mathrm{d}\,\text{Share}_{\text{part}} \approx -v(W_0^{\text{part}})\,\mathrm{d}W_0^{\text{part}}, \quad \mathrm{d}\,\text{Share}_{\text{del}} \approx v(W_0^{\text{del}})\,(-\mathrm{d}W_0^{\text{del}}) - v(W_0^{\text{switch}})\,(-\mathrm{d}W_0^{\text{switch}}),$$

$$\mathrm{d}\,\text{Share}_{\text{self}} \approx -v(W_0^{\text{switch}})\,\mathrm{d}W_0^{\text{switch}},$$

where $v(\cdot)$ is the density of $V(W_0)$.

Formalization via Lorenz ordering. Let $L(p;\theta)$ be the Lorenz curve of $Y(\cdot;\theta)$ and μ the (normalized) mean. For a small perturbation $d\theta$, write the Gateaux variation

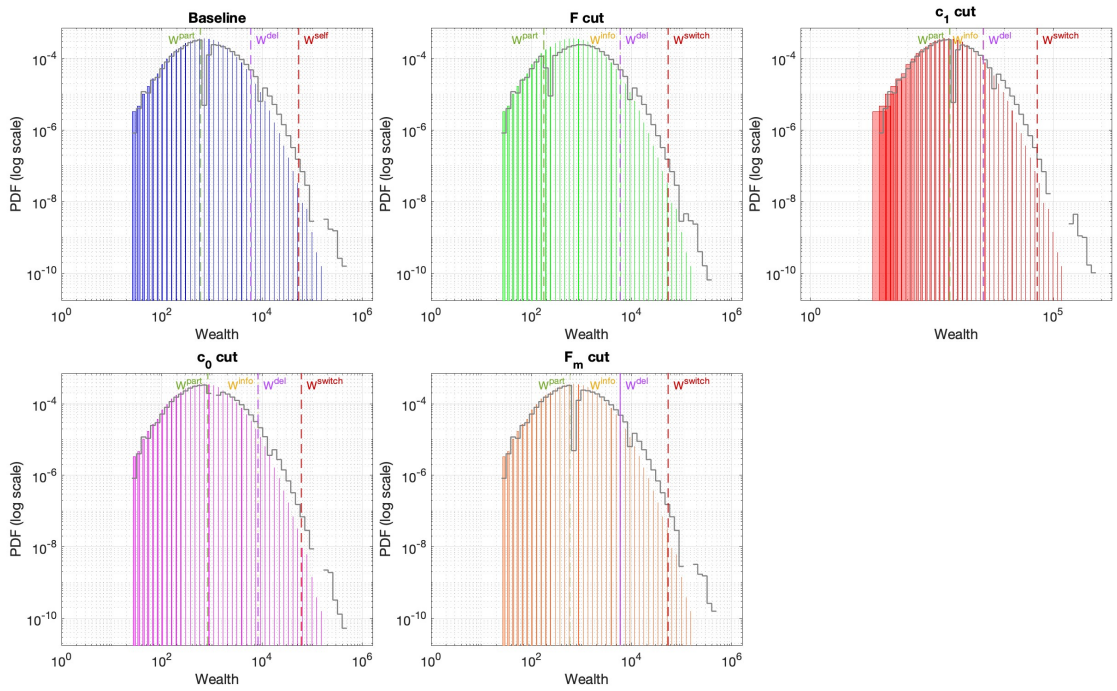
$$dL(p;\theta) = \frac{1}{\mu}\int_0^{q(p)} dY(W;\theta)\,dF(W) - \frac{L(p;\theta)}{\mu}\,d\mu,$$

where $q(p)$ is the pth quantile of F. In the four cases above, $dY(\cdot;\theta)$ is (i) concentrated at the bottom with nonnegative sign (progressive) for $F \downarrow$, (ii) a net increasing function of W on the participants' support for $F_m \downarrow$ under (A1) (regressive), and (iii) an increasing and strictly convex function of W on the participants' support for $c_1 \downarrow$ and $c_0 \downarrow$ (regressive). The Gini is $G(\theta) = 1 - 2\int_0^1 L(p;\theta)\,dp$; thus $dG = -2\int_0^1 dL(p;\theta)\,dp$, which inherits the signs above by the Dalton principle (progressive $\Rightarrow dL \geq 0 \Rightarrow dG \leq 0$, regressive $\Rightarrow dL \leq 0 \Rightarrow dG \geq 0$).

Combining all parts yields the four scenario signs stated in Corollary 1. □

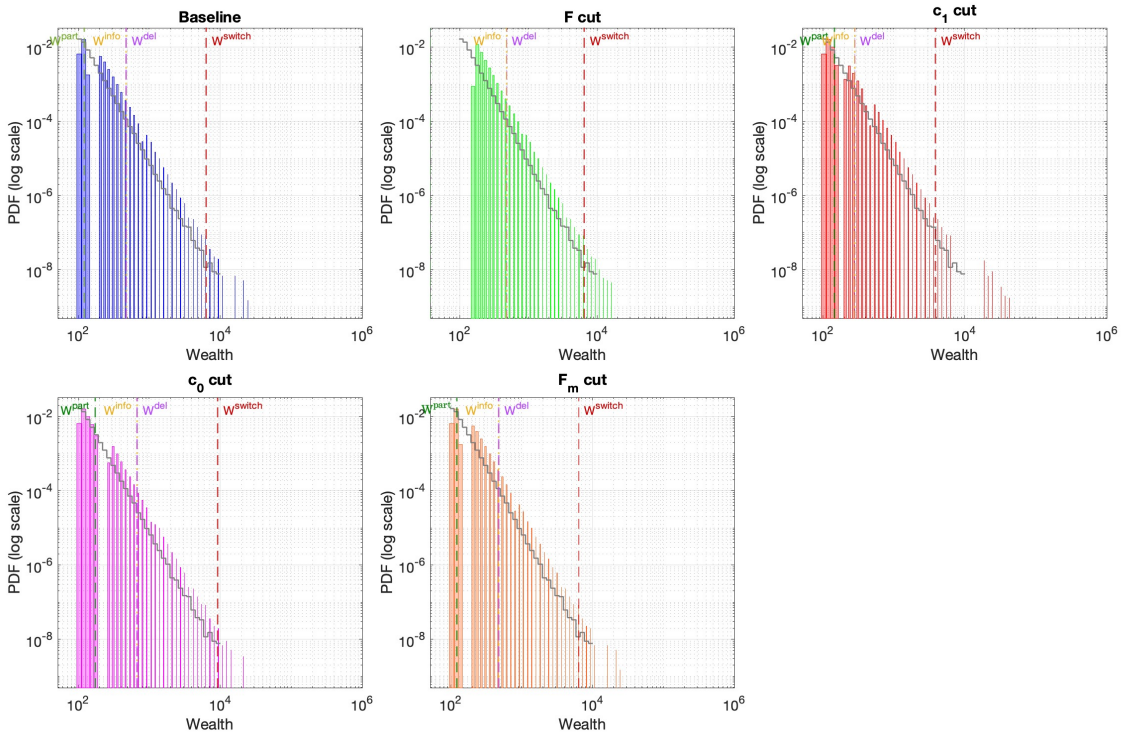
D.9 Additional tables and figures for delegating extension

Figure 16: **Final wealth distributions and endogenous cutoffs (Lognormal fit).** Each panel shows the histogram of terminal wealth W_2 (log scale) for the Baseline (blue) and for one-at-a-time 70% cuts in the participation cost F (green), and 50% cut in the marginal information cost c_1 (red), the quadratic information cost c_0 (magenta), and the search for manager cost F_m (orange). Vertical lines mark the endogenous entry threshold W_{part} (green) and the information threshold W_{info} (orange). Cutting F shifts W_{part} far left while W_{del} moves little, compressing the distribution toward the center; cutting c_1 shifts W_{del} left relative to W_{part}, thickening the right tail; cutting c_0 moves both thresholds right together (their ratio is unchanged), leaving the overall shape of W_2 close to Baseline.



Notes. Baseline: lognormal initial wealth W_0 with parameters held fixed across scenarios; primitives $(\rho, h_0, \sigma_x^2, r_{\text{bond}}) = (0.90, 0.1, 1, 1.07)$ and baseline costs $(F, c_1, c_0) = (2, 2, 2)$ and $F_m = 50$. Scenarios apply a 50% cut to the named cost only, except for the F cut, where we apply a 70% cut; all other primitives are unchanged. For each scenario we solve the fixed-point conditions for (s, t, T_m, f, h_j, h_m), compute W_{part}, W_{del}, W_{info} and W_{switch}, map W_0 into W_2, and plot the resulting histogram using common log-spaced bins; the initial W_0 density is overlaid as a gray outline. Colors match the inequality figure (baseline blue; F cut green; c_1 cut red; c_0 cut magenta, F_m cut orange). Variable definitions are in Appendix A.

Figure 17: **Final wealth distributions and endogenous cutoffs (Pareto fit).** Each panel shows the histogram of terminal wealth W_2 (log scale) for the Baseline (blue) and for one-at-a-time 70% cut in the participation cost F (green), and 50% cut in the marginal information cost c_1 (red), the quadratic information cost c_0 (magenta), and the search for manager cost F_m (orange). Vertical lines mark the endogenous entry threshold W_{part} (green) and the information threshold W_{info} (orange). Cutting F shifts W_{part} far left while W_{del} moves little, compressing the distribution toward the center; cutting c_1 shifts W_{del} left relative to W_{part}, thickening the right tail; cutting c_0 moves both thresholds right together (their ratio is unchanged), leaving the overall shape of W_2 close to Baseline.



Notes. Baseline: Pareto initial wealth W_0 with parameters held fixed across scenarios; primitives $(\rho, h_0, \sigma_x^2, r_{\text{bond}}) = (0.90, 0.1, 1, 1.07)$ and baseline costs $(F, c_1, c_0) = (2, 2, 2)$ and $F_m = 50$. Scenarios apply a 50% cut to the named cost only, except for the F cut, where we apply a 70% cut; all other primitives are unchanged. For each scenario we solve the fixed-point conditions for (s, t, T_m, f, h_j, h_m), compute W_{part}, W_{del}, W_{info} and W_{switch}, map W_0 into W_2, and plot the resulting histogram using common log-spaced bins; the initial W_0 density is overlaid as a gray outline. Colors match the inequality figure (baseline blue; F cut green; c_1 cut red; c_0 cut magenta, F_m cut orange). Variable definitions are in Appendix A.