Financial Technology and Inequality

Roxana Mihet[‡]
Swiss Finance Institute at HEC Lausanne

January 27, 2026

Abstract

Financial technology that improves participation in equity markets can also favor sophisticated investors. I show this promise-peril tension in a noisy rational-expectations portfolio choice model with investor demand heterogeneity, by studying the distributional incidence of three distinct technology shocks to entry, information adoption, and learning precision. Lower entry costs broaden risk sharing and reduce wealth and return inequality. Cheaper information reallocates rents toward informed investors and increases cross-sectional dispersion. Reductions in precision costs amplify inequality within the informed group. These implications are robust to heterogeneity in absolute risk-aversion and to the possibility that households delegate information acquisition to an intermediary.

Keywords: Financial technology, stock market, inequality, Gini coefficient, participation, information frictions, wealth distribution.
JEL Classification: E21, G11, G14, L1, L15.

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

†The author is grateful to Anna Bayona, Jess Benhabib, Stijn Claessens, Jerome Dugast, Jon Frost, Andreas Fuster, Leonardo Gambacorta, Diego Garcia, Mark Gertler, Avi Goldfarb, Amit Goyal, Paymon Khorrami, Luc Laeven, Joseba Martinez, Guido Menzio, Emily Moschini, John Muellbauer, Cecilia Parlatore, Luigi Pistaferri, Thomas Philippon, Claudia Sahm, Thomas Sargent, Lin Shen, Hyun Shin, Johannes Stroebel, Victoria Vanasco, Laura Veldkamp, Venky Venkateswaran, Xavier Vives, Ansgar Walther and Anthony Lee Zhang, as well as participants at numerous conferences and workshops for their helpful comments and discussions. This paper is based on a chapter of the author's dissertation at NYU Stern, where it won the NYU Best Dissertation Award in the Inequality Category, the ECB Young Economist Best Paper Award, the EEA Best Paper Special Mention Award, and the WFA Cubist Strategies Prize for Outstanding Research. The author acknowledges generous funding from the Sandoz Foundation - Monique de Meuron Programme for Academic Excellence, the Internet Association Grant, and the BFI Macro-Financial Modeling Dissertation Grant. All errors are the author's responsibility.

1 Introduction

Financial technology is changing both *who* can access capital markets and *who* can extract rents from trading in them. Brokerage platforms and automation have lowered onboarding and transaction costs, while data, machine learning, and generative AI have lowered the private cost of producing and exploiting signals. Yet the household-finance facts that motivated the modern literature remain: stock market participation is incomplete and heterogeneous, and realized portfolio returns are highly dispersed across households, with large differences between sophisticated and unsophisticated investors.[1] I provide a theory of why these trends can coexist in the presence of financial innovation: *technological progress that broadens access is not the same as technological progress that sharpens informational advantages.* This access–versus–advantage distinction is central to the promise–peril view of technology in personal finance (Campbell and Ramadorai, 2025): the distributional effects of financial technology hinge on whether it lowers the cost of entering markets or the cost of producing and exploiting information.

I develop a parsimonious equilibrium model of investing with two workhorse frictions in household finance: a fixed cost of equity-market participation and a cost of acquiring information, used extensively to explain nonparticipation and heterogeneity in risky-asset holdings (see Vissing-Jorgensen, 2002b; Gomes and Michaelides, 2005 for canonical participation-cost portfolio-choice frameworks). The model is a noisy rational expectations economy in which prices aggregate information, but investors can endogenously choose whether to enter the market and whether to acquire a signal. I use the model as a *distributional accounting framework* to characterize the incidence of distinct technology-induced cost reductions on the cross-section of *certainty-equivalent* (CE) returns, and therefore into Lorenz curves, Gini coefficients, and related inequality objects.

[1] See, e.g., Gomes, Haliassos and Ramadorai (2021) for a synthesis of household finance and the participation/return-dispersion facts, and Campbell, Ramadorai and Ranish (2019); Calvet, Bach and Sodini (2019); Fagereng, Guiso, Malacrino and Pistaferri (2020); Smith, Zidar and Zwick (2023) for evidence on dispersion in realized returns and wealth outcomes.

Technology affects three primitives: (i) a fixed participation cost (i.e., onboarding, basic literacy, and administrative frictions), (ii) an *adoption* cost of becoming informed (i.e., set-up and fixed investments required to turn data into actionable signals), and (iii) an *intensity* cost governing how precise a signal is once an investor is informed (i.e., ongoing marginal costs of improving precision, such as data engineering, compute, and expertise).

This separation yields an incidence decomposition that is both simple and empirically interpretable. In the model, inequality moves through two components. The first is *between-group inequality*: differences in CE payoffs between investors who remain uninformed and those who become informed (and, when participation is costly, between participants and nonparticipants). The second is *within-informed inequality*: dispersion *among the informed* driven by heterogeneity in chosen information intensity (precision). This mapping yields sharp theoretical rankings for how technology shifts the between-group and within-informed components of inequality, and it links these movements to realized return dispersion and trading volume.

The main results can be summarized as follows. First, lowering *participation* costs reduces wealth and return inequality by shrinking the mass of nonparticipants and compressing the CE gap between participants and nonparticipants. Second, lowering the *adoption* cost of information raises the between-group component of inequality by shifting investors across a discrete extensive margin: it expands the informed group in a way that reallocates informational rents and increases dispersion between the newly informed and those who remain uninformed. Third, lowering the *intensity* cost primarily operates through the intensive margin: it increases information precision among those who are already informed, amplifying within-informed dispersion. The net effect on aggregate inequality depends on the equilibrium response of participation and prices, and I provide sufficient conditions (and a numerical illustration) for the overall ranking.[2]

[2]The cost shape matters. With a fixed information component, there exists a sharp threshold between being uninformed and informed, so technology that lowers information adoption has a strong

Together, these results imply that improvements in access can coincide with rising cross-sectional dispersion in returns, even as participation becomes more widespread. These incidence results do not require a participation margin; the universal-participation limit preserves the adoption vs. intensity distinction.

The model implies that policy evaluations of FinTech and AI should not treat "lower costs" as a single object. Policies that look similar in the language of "improving access" can operate on different margins and therefore have different incidence. Concretely, reducing onboarding/compliance frictions, account minimums, and operational fixed costs targets the participation wedge and is predicted to raise participation while compressing dispersion in CE payoffs at the bottom and middle of the distribution. By contrast, subsidizing or deregulating proprietary data access, allowing persistent preferential access to datasets, or enabling exclusive data-driven trading advantages primarily targets the information wedge and can *increase* dispersion in realized returns among participants, even if it increases price informativeness. The policy takeaway is therefore *a complementarity, not a slogan*: if regulators and platforms want "more finance for more people" *without* amplifying return inequality, participation-enhancing reforms need to be paired with guardrails that limit the concentration of informational rents (e.g., disclosure and access rules for market-relevant data, constraints on exclusivity that creates persistent informational advantages, and design choices that steer retail innovation toward lowering fixed entry frictions rather than intensifying an arms race in private precision). The framework also suggests a practical diagnostic for policy: track the joint movement of (i) participation and trading volume and (ii) cross-sectional dispersion in realized returns. Innovations that raise (i) while lowering (ii) are likely participation-cost shocks; innovations that raise both are likely information-cost shocks.

Peress (2004, 2005) studies equilibrium participation and endogenous information acquisition in noisy rational-expectations economies, with the goal of explaining mar-

between-group effect. With solely marginal precision costs, everyone learns, and inequality movements operate through within-informed dispersion.

ket evolution—how costly information and heterogeneous demands jointly determine prices, price informativeness, and who participates in equity markets. My paper takes that equilibrium discipline as given, but asks a different question: the distributional incidence of technology shocks. This shift in lens yields new mechanisms for how financial technology impacts inequality, allowing the model to speak to the promise–peril incidence of technology adoption across the wealth distribution.

Finally, I complement the analytical results with a numerical illustration that disciplines the extensive margins and quantifies the relative importance of the extensive (group shares) versus intensive (within-informed) channels under standard wealth priors. The quantitative exercise is not an estimation; it is intended to show that the comparative statics operate at economically meaningful scales and to make the decomposition transparent.

1.1 Related Literature

Household finance: participation and return dispersion. The paper is anchored in the household-finance literature documenting incomplete stock market participation, heterogeneity in risky positions, and large dispersion in realized returns. Gomes et al. (2021) provides a comprehensive synthesis and emphasizes the centrality of participation frictions, heterogeneity in effective risk-bearing capacity, and institutional features that shape household portfolios. Classic evidence on nonparticipation and fixed costs includes Vissing-Jørgensen (2002a); Vissing-Jorgensen (2002b) and related work; the life-cycle implications of participation costs and portfolio choice are developed in Gomes and Michaelides (2005). The paper also connects to evidence documenting the importance of the upper tail and dispersion in realized returns (e.g., Calvet et al. (2019), Campbell et al. (2019), Benhabib, Bisin and Luo (2017), Benhabib, Bisin and Luo (2019)). I build on these insights by embedding participation and information acquisition in an equilibrium setting in which prices respond endogenously,

returns increase with wealth, and by focusing on inequality objects for CE payoffs.

Endogenous information in equilibrium asset markets. The analysis builds on the noisy rational expectations tradition in which prices aggregate dispersed information and investors endogenously acquire signals (e.g., Grossman and Stiglitz, 1980; Verrecchia, 1982; Admati and Pfleiderer, 1990). More broadly, the paper relates to work on endogenous information acquisition and portfolio concentration (e.g., Veldkamp, 2006; Van Nieuwerburgh and Veldkamp, 2009, 2010); my focus differs by isolating adoption vs. precision costs inside a Grossman–Stiglitz trading environment and tracing the incidence on CE inequality objects. Related theory highlights increasing returns to information and wealth-dependent learning choices (e.g., Peress (2004); Kacperczyk, Nosal and Stevens (2018); Azarmsa (2022)), while empirical evidence suggests that informational advantages are concentrated among wealthier households and intermediaries and are associated with higher realized returns (e.g., Frost, Gambacorta and Gambacorta, 2022; Bretscher, Gomes, Sabbatucci and Tamoni, 2022). Complementary evidence from market microstructure shows that changes in trading technology alter retail behavior and the information environment (e.g., Martineau and Zoican (2023)).

Financial innovation, data/AI, and distributional consequences. A growing literature argues that modern financial technology simultaneously lowers some frictions and amplifies others, so improved access need not translate into more equal outcomes. Work on FinTech, market design, and data-driven intermediation emphasizes the role of platforms, data, and algorithmic tools in shaping participation and information production (e.g., Philippon, 2020; Duffie, Foucault, Veldkamp and Vives, 2022; Foucault, Gambacorta, Jiang and Vives, 2025; Grennan and Michaely, 2021; Katona, Painter, Patatoukas and Zeng, 2025). Empirical evidence also stresses that abundant data does not mechanically translate into useful information: extracting value requires organizational capacity and expertise (e.g., Banerjee, Davis and Gondhi, 2018; Dessaint, Fou-

cault and Frésard, 2024; Fedyk, Mihet, Rishabh and Gomes, 2023; Mihet, Rishabh and Gomes, 2026). The framework in this paper formalizes that distinction by separating participation, information adoption, and information intensity as distinct levers and tracing their incidence on inequality.

Inequality, learning, and intermediation. The paper also relates to work on delegation and information production in asset management where scale economies in research and frictions in manager selection shape who becomes informed and by how much (e.g., Berk and Green, 2004; Berk and van Binsbergen, 2016; Gârleanu and Pedersen, 2018; Greig, Ramadorai, Rossi, Utkus and Walther, 2024; Bonelli and Foucault, 2024). The delegation extension in this paper treats intermediation as an implementation of information acquisition and shows that the qualitative incidence results persist when households share the costs of information acquisition by investing through funds.

The rest of the paper is organized as follows. Section 2 presents the model and equilibrium. Section 3 proves the distributional results. Section 4 provides a numerical illustration. Section 5 discusses generalizations and two theoretical extensions. Section 6 concludes.

2 Economy with Participation and Information Frictions

This section develops a parsimonious noisy rational expectations economy designed to study the impact of progress in financial information technology on capital wealth and return inequality. The key modeling choice is to represent technological progress as shocks to three separable primitives: a fixed market-entry cost, an adoption component of information costs (an extensive margin), and an information-intensity component (an intensive margin).

2.1 Environment

Agents. A unit mass of competitive investors $j \in [0,1]$ differ only in initial wealth $W_{0j} \in (0, \bar{W}_0]$, distributed according to V. Initial risky-asset holdings are normalized to zero; all initial wealth is held in the bond at $t = 0$.

Assets and payoffs. There is a risk-free bond in perfectly elastic supply, priced at 1 and paying gross return $r > 0$. There is one risky asset that pays a fundamental dividend z at $t = 2$, with

$$z \sim \mathcal{N}(\mu_z, \sigma_z^2), \qquad h_0 \equiv \sigma_z^{-2}.$$

The risky asset is supplied by noise traders with (random) net supply

$$x = \mu_x + v, \qquad v \sim \mathcal{N}(0, \sigma_x^2),$$

independent of z and of all private signals.

Preferences. At the trading date, investors choose risky holdings to maximize a mean–variance objective that is a standard local approximation to CRRA and delivers decreasing absolute risk aversion:

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

Thus absolute risk tolerance is increasing in wealth: $1/\rho(W_{0j}) = W_{0j}/\rho$.

2.2 Participation and information technology

Participation (entry) cost. To trade the risky asset, an investor must pay a fixed market-entry cost $F > 0$ at $t = 0$. Let $\mathbf{1}_{\text{part},j} \in \{0,1\}$ denote participation. If $\mathbf{1}_{\text{part},j} = 0$, the investor holds only the bond.

Private information and its cost. A participating investor may acquire a private signal about z of the form

$$s_j = z + \varepsilon_j, \qquad \varepsilon_j \sim \mathcal{N}(0, \sigma_{s,j}^2), \qquad h_j \equiv \sigma_{s,j}^{-2} \geq 0,$$

where h_j is the chosen precision. Acquiring precision h_j costs

$$C(h_j) = c_1 h_j + \frac{c_0}{2} h_j^2, \qquad c_1 > 0,\ c_0 > 0. \tag{2}$$

The parameter c_1 captures an *adoption* component (a strictly positive marginal cost at the origin), which generates an extensive margin (uninformed vs. informed). The parameter c_0 governs *intensity* (the convex marginal cost of improving precision once informed).[3]

Budget constraint. Given participation, information choices, and risky holdings q_{1j} chosen at $t = 1$, terminal wealth is

$$W_{2j} = r\Big(W_{0j} - \mathbf{1}_{\text{part},j} F - C(h_j) - p\, q_{1j}\Big) + q_{1j} z. \tag{3}$$

A nonparticipant has $\mathbf{1}_{\text{part},j} = 0$, $q_{1j} = 0$, $h_j = 0$, and $W_{2j} = r W_{0j}$.

2.3 Timing

The economy unfolds over three dates:

$t = 0$ — **Planning.** Before observing prices or private signals, each investor chooses whether to enter ($\mathbf{1}_{\text{part},j}$) and, conditional on entry, chooses private-signal precision h_j.

[3] When $c_1 \neq 0$, there is a positive marginal cost of learning (i.e., the 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 $c_1 = 0$, $W_0^{info} = 0$, which is unrealistic, although mathematically more tractable. Because $c_1 > 0$, the equilibrium needs to be solved through a fixed point.

$t = 1$ — **Trading.** A competitive market maker posts a price p that clears the market. Participants observe p; informed participants also observe s_j. They then choose q_{1j}.

$t = 2$ — **Consumption.** The dividend z is realized and investors consume 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.

2.4 Equilibrium

We focus on linear rational expectations equilibria in which the risky-asset price is affine in (z, x):

$$rp = a + bz - cx, \tag{4}$$

for constants (a, b, c) determined in equilibrium.

Definition 1 (Linear equilibrium)**.** *A linear equilibrium consists of (i) a linear pricing rule (4); (ii) participation and information rules* $(\mathbf{1}_{part,j}, h_j)_{j\in[0,1]}$ *chosen at* $t = 0$*; and (iii) trading rules* $(q_{1j})_{j\in[0,1]}$ *at* $t = 1$*, such that:*

1. *given the price rule, each investor's choices maximize (1) subject to (3) and the information structure;*
2. *markets clear for the risky asset for every realization of* (z, x)*;*
3. *beliefs are Bayesian and consistent with the price rule and the signal structure.*

For ease, all variables are defined in the Online Appendix A, Table 1.

2.5 Equilibrium characterization and sufficient statistics

A convenient feature of CARA/mean–variance with normality is that equilibrium can be summarized by two aggregates that capture *risk-bearing capacity* and *information*

capacity. Define

$$t \equiv \int_{j:\ \mathbf{1}_{\text{part},j}=1} \rho(W_{0j})^{-1}\, dV, \qquad s \equiv \int_{j:\ h_j>0} \rho(W_{0j})^{-1}\, h_j\, dV. \tag{5}$$

The scalar t is aggregate risk tolerance of participants; the scalar s is aggregate risk-tolerance-weighted private precision of the informed sector.

Proposition 1 (Linear equilibrium: demands, information, and cutoffs)**.** *In any linear equilibrium, the following hold:*

(i) Trading demand. *Given beliefs at $t = 1$, a participant's optimal risky demand is*

$$q_{1j} = \rho(W_{0j})^{-1}\, \frac{\mathbb{E}_j[z] - rp}{\mathbb{V}_j[z]}. \tag{6}$$

(ii) Price coefficients and price informativeness. *Market clearing implies a linear price rule* (4) *whose coefficients are functions of (s, t). In particular, the information content of prices can be summarized by an endogenous* price precision

$$h_p = \frac{s^2}{\sigma_x^2}, \qquad \bar{h} \equiv h_0 + h_p, \tag{7}$$

so that the price aggregates information more strongly when the informed sector's information capacity s is larger.

(iii) Marginal value of private precision. *At $t = 0$, the (ex-ante) marginal mean–variance gain from one additional unit of private precision can be written as*

$$D(s,t) \equiv \frac{\sigma_x^2 + 2st + t^2 h_0}{t^2\,(h_0 + s^2/\sigma_x^2)^2}. \tag{8}$$

For an investor with wealth W_{0j} (risk tolerance $\rho(W_{0j})^{-1}$), the marginal value of precision scales with $\rho(W_{0j})^{-1}$.

(iv) Optimal information intensity. *Conditional on choosing to become informed, the optimal precision schedule is linear in wealth:*

$$h_j(W_{0j}) = \max\left\{0,\ \frac{W_{0j}D(s,t) - 2\rho c_1}{2\rho c_0}\right\}. \tag{9}$$

(v) Cutoffs and sorting. *Participation and information acquisition are monotone in wealth and characterized by cutoffs*

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

Hence $W_0^{info}/W_0^{part} = c_1/F$. *When* $c_1 > F$, *the population partitions into three groups: nonparticipants* ($W_0 < W_0^{part}$), *uninformed participants* ($W_0^{part} \le W_0 < W_0^{info}$), *and informed participants* ($W_0 \ge W_0^{info}$).

(vi) Fixed point. *Given* (s,t), *equations* (9)–(11) *pin down the induced participation and information rules, which in turn imply* (s,t) *via* (5). *Thus equilibrium is characterized by a two-dimensional fixed point in* (s,t).

Proofs and the explicit expressions for (a,b,c) are provided in Appendix B.1-B.7.

Interpretation. The sufficient statistics (s,t) have a direct economic meaning: t is aggregate risk-bearing capacity, while s is the informed sector's information capacity weighted by risk-bearing. The object $D(s,t)$ is the marginal value of private precision at the planning stage. It falls when the market becomes more informative (higher s increases h_p) and when risk-bearing expands (higher t compresses risk premia), yielding clean comparative statics once its slopes are signed on an interior region (formalized in Appendix B.4).

Remark (learning technology adoption vs. learning intensity and why $c_1 > 0$ matters). If $c_1 = 0$ (purely quadratic costs), the model collapses the extensive information margin: all participants choose strictly positive precision. The adoption

component $c_1 > 0$ creates a nontrivial uninformed–informed boundary W_0^{info} and is therefore essential for separating the incidence of technology that lowers the fixed cost of becoming informed from technology that lowers the marginal cost of improving precision once informed.

3 Progress in financial information technology and inequality

This section isolates the equilibrium objects through which financial technology affects the cross-section of outcomes and then maps these objects into standard measures of inequality, such as Lorenz-curves and Gini coefficients. In the model, technology is a vector of cost shifters $\theta \equiv (F, c_1, c_0)$, where F is a fixed participation cost, c_1 is an *adoption* cost of becoming informed, and c_0 governs the *intensity* (curvature) cost of information precision.

3.1 Technology shocks as distributional levers

Two endogenous cutoffs as sufficient statistics. The distributional incidence of θ operates through two endogenous wealth thresholds: the participation cutoff W_0^{part} and (conditional on participation) the information-adoption cutoff W_0^{info}. In equilibrium,

$$W_0^{\text{part}} = \frac{2\rho F}{D(s,t)}, \qquad W_0^{\text{info}} = \frac{2\rho c_1}{D(s,t)}, \tag{11}$$

where s is aggregate private precision, t is aggregate risk tolerance, and $D(s,t)$ is the marginal value (in utility units) of an additional infinitesimal unit of private precision. Only sign properties of D matter for the comparative statics below.

Throughout, I focus on the empirically relevant *active-information* region $s > \sigma_x^2/(2t)$, under which $D_t < 0$ and $D_s < 0$ (Appendix B.4.1). This region is the one in which private information is produced in equilibrium and is satisfied by the calibrations reported later.

A first immediate implication of (11) is that the *relative position* of the two cutoffs depends only on the ratio of the fixed components of the two frictions:

$$\frac{W_0^{\text{info}}}{W_0^{\text{part}}} = \frac{c_1}{F}. \tag{12}$$

Thus, shocks to F and c_1 change the *distance* between the two cutoffs (and therefore the relative mass of uninformed participants), whereas shocks to c_0 can only rescale both cutoffs through the common denominator $D(s,t)$.

The next lemma records the intermediate comparative statics that will be used in the incidence decomposition and in the Lorenz/Gini results.

Lemma 1 (Technology shocks and cutoff movements)**.** *Fix an interior equilibrium in the active-information region $s > \sigma_x^2/(2t)$. Consider a one-at-a-time reduction in a single component of $\theta = (F, c_1, c_0)$, holding the other two fixed.*

1. ***Entry-cost reduction (fall in F).*** *The participation cutoff falls while the information-adoption cutoff rises: W_0^{part} falls and W_0^{info} increases. Equivalently, the two cutoffs move apart.*

2. ***Adoption-cost reduction (fall in c_1).*** *The participation cutoff rises while the information-adoption cutoff falls: W_0^{part} increases and W_0^{info} falls. Equivalently, the two cutoffs move toward each other.*

3. ***Intensity-cost reduction (fall in c_0).*** *The ratio $W_0^{\text{info}}/W_0^{\text{part}}$ is unchanged (by (12)). Both cutoffs move only through the common denominator $D(s,t)$, i.e., they are rescaled in the same direction (both up or both down).*

Proof. The proof is in Appendix B.9. The intuition for this result follows. First, a reduction in the entry cost F directly increases participation: the cutoff W_0^{part} falls, so marginal households enter. These marginal entrants are poorer and (on average) remain uninformed, which changes the composition of the trading population. In general equilibrium, entry affects the uninformed Sharpe through two channels summarized by $D(s,t)$. First, more participation expands aggregate risk-bearing capacity ($t \uparrow$), which compresses risk premia and reduces the uninformed Sharpe ($D \downarrow$). Second, because the marginal entrants are predominantly uninformed, entry makes prices less informative ($s \downarrow$), which tends to raise the uninformed Sharpe ($D \uparrow$). The key step is that in the active-information region $D_t < 0$ and $D_s < 0$: a larger risk-bearing sector ($t \uparrow$) and a more informative price ($s \uparrow$) both compress the uninformed Sharpe ($D \downarrow$). A fall in F raises the incentive to participate, but reduces the return to information acquisition: informed households face a lower equilibrium risk premium, so fewer of them find it

worthwhile to pay the information fixed cost on the margin. Hence W_0^{part} falls, while W_0^{info} increases.

Second, the direct effect of lowering the adoption cost c_1 is to make it privately optimal to become informed at lower wealth levels, so the information-adoption cutoff falls ($W_0^{\text{info}} \downarrow$). This expands the informed sector and increases price informativeness ($s \uparrow$). A more informative price reduces the incremental benefit of trading without information: the equilibrium uninformed Sharpe $D(s,t)$ falls when s rises ($D_s < 0$). As the uninformed Sharpe compresses, marginal uninformed traders exit due to adverse selection, unless they are sufficiently wealthy to bear risk at the now-lower reward-to-risk. Thus the participation cutoff rises ($W_0^{\text{part}} \uparrow$). In short, cheaper adoption increases the mass of informed traders, makes prices more revealing, and thereby erodes the rents from uninformed participation in equilibrium through increased adverse selection. Hence W_0^{part} increases, while W_0^{info} falls.

Third, a lower c_0 shifts the *intensive* margin of information: for a given set of informed households, a lower c_0 increases equilibrium signal precision and thus price informativeness ($s \uparrow$). Both wealth cutoffs share the same denominator $D(s,t)$ and differ only in their numerators (F versus c_1), so changes in c_0 primarily act like a *common scaling* of the two thresholds: $W_0^{\text{info}}/W_0^{\text{part}} = c_1/F$ remains fixed. The key step is again that $D_s < 0$ and $D_t < 0$ in the active-information region: movements in (s,t) translate monotonically into movements in the uninformed Sharpe $D(s,t)$. Thus c_0 cannot "rotate" participation versus adoption the way F or c_1 can; it only shifts both cutoffs together through the common general-equilibrium effect on $D(s,t)$.

3.2 Distributional incidence: from technology to inequality

Section 3.1 characterized technology shocks $\theta = (F, c_1, c_0)$ as *distributional levers*. In particular, Lemma 1 shows how changes in θ shift the participation cutoff $W_0^{\text{part}}(\theta)$ and the information-adoption cutoff $W_0^{\text{info}}(\theta)$, thereby reallocating mass across non-

participants, uninformed participants, and informed participants. The remaining step is to map these cutoff movements—together with the endogenous precision dispersion among the informed—into inequality objects. I do so by working with the certainty-equivalent terminal wealth schedule $Y(W_0;\theta)$ and by deriving Lorenz- and Gini-type comparisons from the shape of the induced perturbation $dY(W_0;\theta)$. This delivers an incidence decomposition of financial innovation into between-group and within-informed components.

Fix a technology vector $\theta \equiv (F, c_1, c_0)$ and let $(s(\theta), t(\theta))$ denote the equilibrium aggregates, where s is the informed sector's information capacity weighted by risk-bearing and t is the aggregate risk-bearing capacity, with associated $D(\theta) \equiv D(s(\theta), t(\theta))$ (the market-wide expected squared Sharpe ratio for an uninformed participant, equivalently the marginal value of an infinitesimal private signal). The distributional analysis in this section treats equilibrium as a mapping from initial wealth W_0 into a *certainty-equivalent* terminal wealth schedule. Because preferences are CARA and payoffs are normal, the investor's planning-date problem admits a mean–variance representation, so distributional comparisons can be stated directly in terms of certainty-equivalent (risk-adjusted) wealth.

A certainty-equivalent wealth schedule and its two wedges. Define the certainty-equivalent terminal wealth of an investor with initial wealth W_0 as

$$Y(W_0;\theta) \;\equiv\; rW_0 \;+\; \underbrace{\Big(\frac{W_0}{2\rho}D(\theta) - F\Big)_+}_{\text{market-entry wedge}} \;+\; \underbrace{\Big(\frac{\big(W_0 D(\theta) - 2\rho c_1\big)^2}{8\rho^2 c_0}\Big)_+}_{\text{information wedge}}. \tag{13}$$

The first term is the common safe accumulation of initial wealth. The second term is the net certainty-equivalent surplus from *entering* the market without private information (the entry wedge). The third term is the additional certainty-equivalent surplus from *becoming informed* and optimally choosing signal precision given the quadratic information technology (the information wedge). The "$\cdot_+$" operator makes explicit that

these wedges activate only above endogenous cutoffs.

Equation (13) makes the two primitive components of inequality clear:[4]

- *Between-group inequality* is governed by the activation of the wedges: below the participation cutoff investors earn only the outside option rW_0, while above it they earn the entry wedge; above the information cutoff they additionally earn the information wedge.

- *Within-informed dispersion* is governed by the curvature of the information wedge: conditional on being informed, the information term is strictly convex in W_0 (through the optimal scaling of precision and positions), creating dispersion within the informed group even holding group shares fixed.

The corresponding cutoffs are

$$W_0^{\text{part}}(\theta) = \frac{2\rho F}{D(\theta)}, \qquad W_0^{\text{info}}(\theta) = \frac{2\rho c_1}{D(\theta)}, \tag{14}$$

and for $W_0 \geq W_0^{\text{info}}(\theta)$ the optimal private precision (in the notation of the model) is linear in wealth,

$$\sigma_s^{-2}(W_0;\theta) = \frac{W_0 D(\theta) - 2\rho c_1}{2\rho c_0}\,. \tag{15}$$

This wealth-linearity is the formal sense in which informed investors optimally "time and size": higher W_0 raises risk tolerance and therefore amplifies the optimal intensity of information acquisition and the scale of positions taken conditional on signals.

From a wealth schedule to inequality objects. Let V denote the initial-wealth distribution, such that $W_0 \sim V$. Equation (13) induces a distribution over certainty-equivalent terminal wealth, $Y_\theta \equiv Y(W_0;\theta)$, and hence Lorenz curves, Gini coefficients,

[4]The adoption term c_1 pins the marginal cost at zero precision and therefore governs the *extensive* decision of whether to become informed at all: it determines who crosses the uninformed/informed boundary. By contrast, c_0 governs the curvature of information costs and therefore primarily governs the *intensive* precision choice conditional on being informed. This separation is exactly what allows a clean decomposition of inequality movements into (i) between-group effects driven by cutoff shifts, and (ii) within-informed effects driven by precision dispersion.

and top-share statistics. The comparative statics results below are therefore stated as properties of the mapping $\theta \mapsto Y(\cdot;\theta)$: technology matters only through how it shifts the entry wedge, the information wedge, and the endogenous cutoffs in (14), holding fixed the primitive V.

Let $\mu(\theta) \equiv \mathbb{E}[Y_\theta]$, and let $Q_Y(p;\theta)$ denote the quantile function of Y_θ. The Lorenz curve is

$$L(p;\theta) \equiv \frac{1}{\mu(\theta)} \int_0^p Q_Y(u;\theta)\, du, \qquad p \in [0,1],$$

and the associated Gini coefficient is

$$G(\theta) \equiv 1 - 2\int_0^1 L(p;\theta)\, dp.$$

For tail comparisons I also report top-share statistics

$$S_\alpha(\theta) \equiv \frac{1}{\mu(\theta)} \int_{1-\alpha}^1 Q_Y(u;\theta)\, du, \qquad \alpha \in (0,1).$$

To isolate within-informed dispersion, define the conditional distribution of Y_θ given $W_0 \geq W_0^{\text{info}}(\theta)$, and denote its Lorenz curve and Gini coefficient by $L^{\text{info}}(\cdot;\theta)$ and $G^{\text{info}}(\theta)$.

3.3 Incidence results: technology shocks and inequality

Theorem 1 states the main distributional incidence ranking and connects it to the cutoff comparative statics in Lemma 1.

Theorem 1 (Incidence decomposition and inequality rankings)**.** *Fix an interior equilibrium in the active-information region and consider one-at-a-time local reductions in the technology vector $\theta = (F, c_1, c_0)$.*

1. ***Participation technology (fall in F).*** *A reduction in the entry cost F is Lorenz-improving: $L(p;\theta)$ shifts weakly up for all p and the Gini coefficient $G(\theta)$*

falls. The mechanism is a progressive reallocation across the participation cutoff (Lemma 1) that compresses between-group gaps.

2. ***Information-adoption technology (fall in*** c_1***).*** *A reduction in the information-adoption cost* c_1 *is Lorenz-worsening:* $L(p;\theta)$ *shifts weakly down and* $G(\theta)$ *rises. The mechanism is a regressive reallocation across the uninformed–informed boundary that thickens the upper tail among participants, raising the between-group component of inequality.*

3. ***Information-intensity technology (fall in*** c_0***).*** *A reduction in the information-intensity cost* c_0 *increases within-informed dispersion:* $L^{\text{info}}(p;\theta)$ *shifts down and* $G^{\text{info}}(\theta)$ *rises. The sign for the* aggregate *Lorenz/Gini is in general ambiguous because the induced increase in price informativeness changes* $D(\theta)$ *and feeds back into the cutoffs and the entry wedge.*

Proof. The proof is in Appendix C and uses the Gateaux variation of the Lorenz curve and the Dalton principle to sign changes in $L(p;\theta)$ and $G(\theta)$ from the shape of the perturbation $dY(W_0;\theta)$ induced by each technology shock.

Economic intuition (through group Sharpe ratios). In this CARA–Normal economy, equilibrium investment opportunities admit a sharp summary: for any trading-date information set $\mathcal{I}_j$, the planning-date certainty-equivalent gain from optimally taking risk equals

$$\frac{\tau(W_0)}{2}\,\mathbb{E}\big[S_j^2(\mathcal{I}_j)\big]\,, \qquad \tau(W_0) \equiv \frac{W_0}{\rho},$$

where $S_j(\mathcal{I}_j)$ is the conditional Sharpe ratio of the risky asset perceived at $t = 1$. Technology therefore affects inequality through two channels: (i) it shifts *who* has access to positive Sharpe opportunities (cutoff movements), and (ii) it shifts the *dispersion* of Sharpe opportunities among those who are informed (precision dispersion).

Mechanism: Sharpe ratios by group. There are three groups. Nonparticipants have no exposure to the risky asset and hence no excess-return opportunity: $S_N^2 \equiv 0$. Uninformed participants condition only on the equilibrium price; their expected *squared* Sharpe ratio is

$$\mathbb{E}[S_U^2] = D(s,t),$$

so $D(s,t)$ is both the marginal value of an infinitesimal private signal and the (risk-adjusted) investment opportunity available to a price-only trader (see Appendix B.10). On the active-information region, $D_t < 0$ and $D_s < 0$: a larger risk-bearing sector ($t \uparrow$) and a more informative price ($s \uparrow$) compress the uninformed Sharpe.

Informed participants observe a private signal and choose its precision. Under quadratic costs, their optimal signal precision is

$$\sigma_{s,j}^{-2}(W_0) = \frac{W_0 D(s,t) - 2\rho c_1}{2\rho c_0} \cdot \mathbf{1}\{W_0 \geq W_0^{\text{info}}\},$$

so the informed group's Sharpe advantage is increasing in $\sigma_{s,j}^{-2}$ and therefore increasing in wealth (and more sensitive to wealth when c_0 is smaller). This is the within-informed dispersion channel.

Putting the pieces together. Each technology shock changes both (a) *membership* in the three groups (via Lemma 1) and (b) *Sharpe opportunities* within and across groups (via $D(s,t)$ and $\sigma_{s,j}^{-2}$):

1. *Participation technology (fall in F).* Lower F expands participation at the bottom: W_0^{part} falls. The marginal entrants are predominantly uninformed, which raises aggregate risk tolerance t and (locally) weakens incentives to acquire information, tending to reduce $D(s,t)$. Thus the uninformed Sharpe opportunity $\mathbb{E}[S_U^2] = D(s,t)$ falls, and the marginal value of private precision falls as well, so information adoption becomes less attractive and W_0^{info} rises. In Sharpe-ratio terms, the economy reallocates mass from a zero-Sharpe outside option into a larger pool

earning a more compressed (lower) Sharpe, while simultaneously shrinking the high-Sharpe informed sector. This is a progressive perturbation of the CE wealth schedule, yielding Lorenz improvement and a lower Gini coefficient.

2. *Information-adoption technology (fall in c_1).* Lower c_1 lowers the information-adoption threshold W_0^{info} and expands the mass of informed agents. The resulting increase in aggregate precision s makes prices more informative, which lowers $D(s,t)$ and compresses the uninformed Sharpe. At the same time, the direct reduction in c_1 raises the net benefit of becoming informed and shifts the precision schedule up, increasing the Sharpe advantage of informed participants. Hence the Sharpe gap between informed and uninformed widens at the top while uninformed opportunities compress, which is regressive in the Lorenz sense and increases the Gini coefficient.
3. *Information-intensity technology (fall in c_0).* Lower c_0 acts primarily on the intensive margin: conditional on being informed, the mapping from wealth to precision (and thus to Sharpe advantage) becomes steeper, increasing dispersion *within* the informed group. General equilibrium feeds back through an increased s, the informed sector's information capacity weighted by risk-bearing, and lower $D(s,t)$, which compresses Sharpe opportunities for uninformed participants and can shift both cutoffs through the common denominator $D(s,t)$. These offsetting between-group effects make the aggregate Lorenz/Gini impact ambiguous, but the within-informed Lorenz curve shifts down and the within-informed Gini coefficient rises, as stated in Theorem 1.

Robustness/limiting cases These incidence results do not hinge on the presence of a fixed participation cost. Even in the limiting case of universal participation (zero entry cost), reductions in information adoption and information intensity have the same qualitative effects on inequality: adoption primarily shifts mass across the uninformed–

informed boundary (between-group effects), while intensity primarily increases dispersion through precision among the informed (within-informed effects). The participation friction is included because it is empirically first-order and because it delivers additional predictions for participation and trading volume.

4 Quantitative illustration

This section provides a *quantitative illustration* of the incidence forces characterized in Theorem 1. The exercise is deliberately not an estimation: the costs (F, c_1, c_0) are not directly observable and the model is intentionally stylized. The purpose is therefore to (i) gauge whether the incidence effects operate at economically meaningful magnitudes under transparent parameter choices, and (ii) show that the qualitative rankings are robust across standard initial wealth distributions and across a wide range of cost reductions. To keep the paper theory-focused, the main text reports only a small set of baseline counterfactuals and summary figures; the full calibration, solution algorithm, and extensive sensitivity analyses are reported in Appendix D.

Accordingly, I interpret the reported magnitudes as order-of-magnitude elasticities of inequality objects to specific technological cost improvements, and I emphasize internal consistency checks (cutoff movements, group shares, and within-informed dispersion) that map directly into the theory.

4.1 Parameterization and targets

Initial wealth distributions. I treat the initial wealth distribution $W_0 \sim V$ as primitive and report results under three benchmark priors that serve distinct roles. The main text uses a *lognormal* prior, which spreads mass across the middle of the distribution and yields clean, visually interpretable Lorenz-curve movements and Gini coefficients for the baseline incidence figures. To assess the role of the upper tail, Appendix D

repeats the same counterfactuals under a *Pareto* prior, which captures the heavy right tail emphasized in macro-inequality work (e.g., Benhabib et al., 2019; Gomez, 2023) and therefore provides a stress test for mechanisms that load disproportionately on wealthy households. Finally, I also report a *Uniform* prior as a thin-tail benchmark and diagnostic case: by removing tail amplification, it isolates the incidence logic from distributional curvature and confirms that the qualitative ranking of technology shocks is not an artifact of a particular parametric tail shape.

What is targeted. The parameterization disciplines the *technology* parameters (F, c_1, c_0) using a small set of transparent targets that pin down the extensive margins emphasized by the theory:

1. the unconditional risky-asset participation share (baseline 60%),
2. the unconditional informed share (baseline 5%),
3. a mean-preserving normalization of the CE terminal wealth distribution.

Given the equilibrium characterization, the first two targets pin down the locations of the endogenous cutoffs W_0^{part} and W_0^{info} as quantiles of V, while the third removes mechanical level shifts so that inequality comparisons isolate *incidence* rather than average wealth changes. The remaining primitives (preferences, fundamental risk, and noise-trader supply) are held fixed across all counterfactuals; Appendix D reports the exact values and the mapping from targets to parameters.

Mean-preserving normalization. All reported counterfactual distributions are normalized to hold the mean of certainty-equivalent terminal wealth fixed. Formally, after computing the equilibrium mapping $Y(W_0; \theta)$, I apply a common multiplicative rescaling of market payoffs so that $\mathbb{E}[Y(W_0; \theta)]$ is constant across scenarios. This isolates distributional effects from trivial changes in the level of returns.

4.2 Counterfactual technology shocks

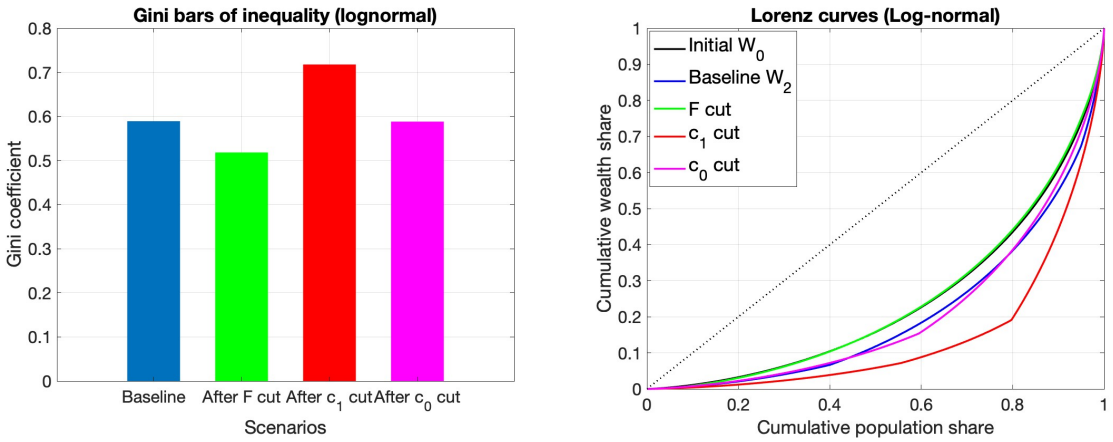
I implement one-at-a-time reductions in each technology cost in $\theta = (F, c_1, c_0)$, holding fixed the return process, the outside option, and the information structure. The benchmark experiments are 70% reductions in F, c_1, and c_0, interpreted as:

1. **Entry technology** ($F \downarrow$): cheaper participation/onboarding,
2. **Information adoption technology** ($c_1 \downarrow$): lower fixed/set-up cost of becoming informed,
3. **Information intensity technology** ($c_0 \downarrow$): cheaper marginal improvements in precision.

The results mirror the theory. Entry technology, shown in green, expands participation at the bottom and compresses between-group gaps, so the Lorenz curve shifts toward equality and the Gini falls to roughly 0.518, a decline of about -0.0710 Gini coefficient points ($\approx -12\%$ relative to baseline). Information adoption technology, shown in red, expands the informed sector and increases the informed advantage relative to uninformed participants, thickening the upper tail and raising the Gini to about 0.71, an increase of 0.129 ($\approx +21\%$ relative). Information intensity technology, shown in magenta, primarily steepens the precision schedule among the informed and thus operates through within-informed dispersion; the aggregate Gini response stays essentially flat and can be ambiguous and dependent on the initial wealth distribution, because general equilibrium improves price informativeness and feeds back into cutoffs and Sharpe ratios.

Figure 1 summarizes these three counterfactuals under the lognormal benchmark, which provides the cleanest visual mapping from each technology lever to Lorenz/Gini movements. Appendix D.7 reports additional figures, including for Pareto (heavy-tail) and Uniform (thin-tail) initial prior distributions, and shows that the qualitative ranking is unchanged across priors; what differs is magnitude: the Pareto prior amplifies

Figure 1: Quantitative illustration under the lognormal prior: baseline and one-at-a-time 70% cost reductions in F, c_1, and c_0 (mean-preserving). Left: Gini coefficient: a higher Gini coefficient corresponds to higher inequality. Right: Lorenz curves: the further the Lorenz curve from the 45-degree line, which denotes perfect equality, the higher wealth inequality is. Appendix D reports full details and additional figures, including Pareto and Uniform priors, and density plots.



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.

the inequality elasticity to information-adoption shocks because more mass sits near or above the information cutoff, while the Uniform prior attenuates tail-driven amplification. Appendix Figure 7 reports the cross-prior comparison of Gini coefficients.

4.3 Decomposition: between-group vs. within-informed

Theorem 1 emphasizes an incidence decomposition: technology affects inequality through (i) between-group reallocation across cutoffs (W_0^{part} and W_0^{info}) and (ii) within-informed dispersion induced by the convexity of the information wedge and the wealth-sensitivity of precision. The quantitative illustration makes this decomposition concrete using two diagnostics.

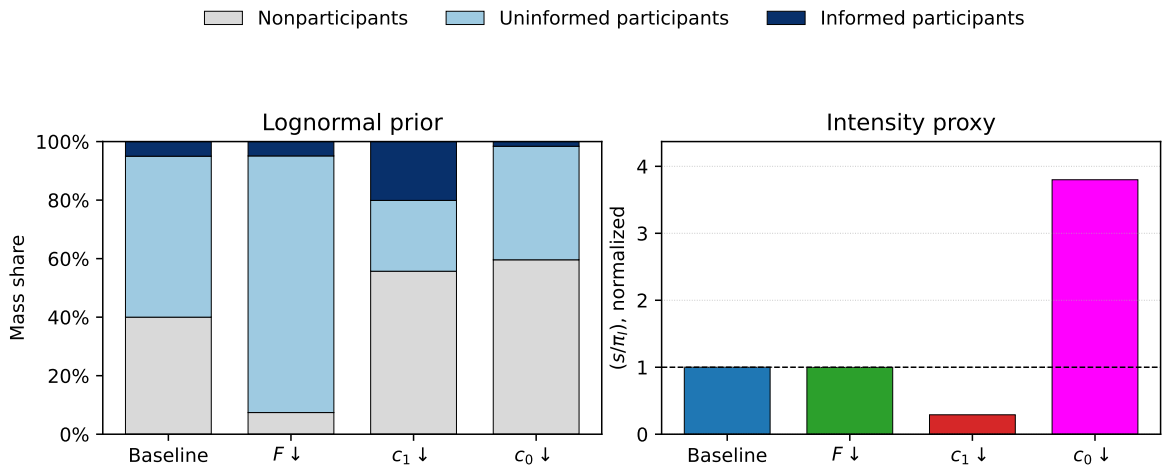
Between-group movements (masses and mean wedges). I track the endogenous group shares and the corresponding group-average CE gains from the entry and information wedges. Entry technology ($F \downarrow$) primarily reduces π_N and compresses the gap between nonparticipants and (mostly) uninformed participants; adoption technology ($c_1 \downarrow$) primarily increases π_I and widens the informed–uninformed gap within participants.

$$\pi_N(\theta) = \Pr(W_0 < W_0^{\text{part}}(\theta)),$$
$$\pi_U(\theta) = \Pr(W_0^{\text{part}}(\theta) \leq W_0 < W_0^{\text{info}}(\theta)),$$
$$\pi_I(\theta) = \Pr(W_0 \geq W_0^{\text{info}}(\theta)),$$

Within-informed dispersion. To isolate the intensive-margin channel, I report inequality within the informed subpopulation, using the conditional Lorenz/Gini objects $L^{\text{info}}(\cdot;\theta)$ and $G^{\text{info}}(\theta)$ introduced in Section 3.2. Intensity technology ($c_0 \downarrow$) produces the clearest movement in these conditional objects: it steepens the mapping from wealth into optimal precision and therefore increases dispersion *within* the informed group, even when aggregate inequality changes little.

Figure 2 summarizes these decomposition diagnostics for the benchmark lognormal illustration; the full set of decomposition tables (including Pareto priors and alternative shock sizes) is reported in Appendix D.7. The left panel isolates the extensive-margin reallocations across groups (nonparticipants, uninformed participants, informed), i.e., the between-group component of inequality. A reduction in the participation cost F primarily shifts mass from nonparticipation into uninformed participation, while a reduction in the adoption cost c_1 shifts mass from uninformed to informed, but also from the uninformed to nonparticipants through the adverse selection channel emphasized above. The right panel reports an intensive-margin diagnostic, (s/π_I), which measures information production per informed investor (normalized to baseline). Consistent with the model's separation between adoption and intensity, lower information

Figure 2: **Quantitative decomposition: extensive vs. intensive margins under the lognormal prior.** Entry technology acts primarily through between-group reallocation; intensity technology acts primarily through within-informed dispersion.



Notes. Left panels report the endogenous mass shares of nonparticipants, uninformed participants, and informed participants across the baseline and one-at-a-time 70% cost cuts. Right panels report an intensity diagnostic, (s/π_I), where s is aggregate information production and π_I is the informed mass share; the statistic is normalized to the baseline within each prior. The right-panel statistic is a diagnostic for the intensive margin (information produced per informed investor), not a separate target. Appendix D reports full details and additional figures, including Pareto priors, and other plots.

adoption costs ($c_1 \downarrow$) expand the informed group, but can dilute average information intensity, whereas lower information intensity costs ($c_0 \downarrow$) sharply raise intensity among the remaining informed.

To make the distributional incidence visually transparent, Appendix Figure 9 overlays the initial wealth distribution (Lognormal, Pareto, and Uniform) and the induced terminal-wealth distributions under baseline costs and under a single cost reduction, separately for entry technology ($F \downarrow$), information adoption technology ($c_1 \downarrow$), and information intensity technology ($c_0 \downarrow$).

4.4 Robustness of numerical exercise

The qualitative ranking is robust across a variety of numerical exercises. Appendix D.7 reports robustness to alternative baseline participation and informed shares, alternative

magnitudes of cost reductions, and additional distributional summaries (including top-share statistics).

Across different priors (Lognormal, Pareto, and Uniform), and across a wide range of one-at-a-time cost reductions, entry technology ($F \downarrow$) consistently reduces inequality while information adoption technology ($c_1 \downarrow$) consistently increases it. In contrast, information intensity technology ($c_0 \downarrow$) consistently increases within-informed dispersion, but has only modest (and potentially offsetting) effects on aggregate inequality, consistent with Theorem 1.

High-participation baseline ("F near zero"). Because F may already be small in modern retail markets, Appendix D.7 repeats the counterfactuals under high baseline participation targets (e.g., 90% and 99%). Consistent with the concern that F reductions become second-order when entry is already cheap, the quantitative effect of $F \downarrow$ on aggregate inequality shrinks in these baselines. In contrast, the adoption and intensity channels remain active: $c_1 \downarrow$ continues to reallocate mass across the uninformed–informed boundary and $c_0 \downarrow$ continues to raise within-informed dispersion.

5 Theoretical generalizations, robustness, and extensions

5.1 Generalizations

The inequality results in the benchmark model do not rely on any single, knife-edge modeling assumption. What the mechanism needs is simply this: *some investors want to hold larger risky stock positions than others* (i.e., heterogeneity in wealth, equivalently in absolute risk tolerance).[5] When that is true, the *value of participating uninformedly* and the *value of becoming informed* are not the same across households, so technology

[5]Many real-world differences generate the required heterogeneity in risky demand. For example, differences in information-processing costs, age and investment horizon, or exposure to background risk all change how attractive stocks are relative to bonds. As long as the heterogeneity affects optimal position size, the paper's participation-versus-learning decomposition continues to apply.

shocks have nontrivial distributional incidence.

Allowing for a participation friction F, as in Peress, 2005, makes this incidence transparent by generating three endogenous states— nonparticipants, uninformed participants, and informed participants—and thus a clean separation between (i) reallocation across the participation cutoff and (ii) reallocation across the uninformed–informed boundary. However, the key inequality forces from *information* technology do *not* require a nonparticipation group. If participation is exogenous or essentially universal (e.g., $F \approx 0$), the economy collapses to two groups (uninformed vs. informed), and the same incidence logic continues to operate: lowering the adoption cost c_1 shifts mass across the uninformed–informed boundary (a between-group channel), while lowering the intensity cost c_0 steepens the wealth-to-precision mapping and increases dispersion *within* the informed group (an intensive channel). The participation margin is therefore needed to speak about entry technology ($F \downarrow$), but it is not needed for the paper's main message that "democratization" of information can coexist with persistent or rising inequality.

CARA versus DARA (including CRRA). It is useful to separate two concepts: (i) whether wealth matters for risky demand, and (ii) whether the model remains tractable.

With *CARA* utility, the number of shares an investor wants to hold does not depend on their wealth (absent borrowing constraints). As a result, wealth by itself does not generate sorting into participation or information acquisition. This is a clean benchmark, but it clashes with basic empirical facts: stock market participation and the stock share both rise strongly with wealth and income (e.g., Lusardi, Michaud and Mitchell, 2017).

With *CRRA* utility, wealth naturally matters for portfolio choice, which is empirically appealing. The drawback is that in a noisy rational expectations setting, CRRA

generally destroys the convenient linearity of equilibrium prices in payoffs and supply, and closed-form equilibrium expressions are typically unavailable. For this reason, I use a tractable approximation that keeps the key empirical feature (richer households behave as if they are more willing to take risk): I adopt a CARA representation evaluated locally around wealth, which implies *decreasing absolute risk aversion*.

What survives under CARA. Even under CARA preferences, the core incidence forces from improved financial information technologies are present *as long as* the model features (i) a participation margin and (ii) coexistence of the three groups (out, uninformed, informed). Lower information costs still increase the advantage of being informed relative to trading uninformed. Informed investors expand their positions and earn higher risk-adjusted returns, while uninformed participants face a worse risk–return tradeoff because prices become more informative.

The main qualitative difference is that, under CARA, a fixed participation cost does *not* create an additional "information channel" through wealth-dependent risk-bearing capacity. In a CARA economy with common absolute risk aversion, the private value of precision does not move with the cross-sectional composition of wealth, so changes in entry costs do not change information acquisition through that route. In particular, in this CARA benchmark the variance of returns monotonically falls when entry costs fall (because participation rises without changing the private incentives to acquire precision in the same way as in the DARA setting).

Timing of information and "early resolution." The model uses a mean–variance structure and a standard timing in which information is produced and traded upon before payoffs are realized. In this environment, better information means that more risk is resolved in the trading period rather than in the consumption period, which increases gains from trade when there is disagreement and heterogeneous exposure.

If, instead, traders could fully trade *before* private information is observed, then

ex-ante payoffs would be linear in wealth and traders would behave effectively risk-neutral at that stage. In that limiting case, the risk premium would collapse to zero. I deliberately abstract from that timing to keep the focus on the paper's main object: how technology-driven changes in the costs of entering and learning reallocate participation and information acquisition across investors.

A dynamic interpretation. The static model maps naturally into a multi-period environment, and it highlights two distinct forces.

First, along the *time-series* dimension, better information technologies shift *when* news gets incorporated into prices. Earlier incorporation can increase price volatility (prices react more strongly to news when it arrives), while reducing the variance of subsequent returns because uncertainty is resolved sooner and risks farther in the future are discounted more heavily (Campbell, Lettau, Malkiel and Xu, 2001). In this sense, the static model's "sharper prices but lower return variance" logic carries over.

Second, along the *cross-sectional* dimension, even small per-period return advantages for informed investors compound over time. Informed investors earn higher Sharpe ratios, uninformed investors lag behind, and these persistent differences generate widening dispersion in wealth trajectories. Dynamic noisy rational expectations models such as Azarmsa (2022) formalize this compounding channel. Empirically, heterogeneous household returns (e.g., Xavier, 2021; Fagereng et al., 2020; Campbell et al., 2019) are consistent with such mechanisms. Appendix E discusses evidence on participation and information frictions and relates the model's mechanisms to long-run trends in wealth dispersion.

Finally, recent evidence by Dessaint et al., 2024 underscores a useful distinction between short-horizon and long-horizon information technologies. Technologies that improve short-run signals can shift volatility forward and reduce short-run return variance, while technologies that improve long-run fundamental information can dispro-

portionately benefit investors who can pay fixed data and processing costs, generating advantages that compound over time. This distinction aligns closely with the model's separation between (i) short-run price informativeness and return variance and (ii) long-run cross-sectional dispersion driven by differential learning.

5.2 Extension 1: Heterogeneity in risk aversion

The model's results are not about "wealth" in isolation. They are about heterogeneity in *effective risk tolerance*, meaning heterogeneity in how large a risky position different investors optimally want to take. Investors who optimally take larger risky positions benefit more from becoming informed (because information is applied to a larger position), so they are more willing to pay the fixed and variable costs of learning. This creates sorting into participation and information acquisition, which is the backbone of the paper's inequality results.

In the baseline model, effective risk tolerance rises with wealth because I use a DARA (CARA-approximation-to-CRRA) structure. In particular, absolute risk aversion falls with wealth, which implies that richer households behave as if they have higher risk tolerance and therefore take larger stock positions, value information more, and choose higher precision.

Why wealth heterogeneity alone is not sufficient under CARA. To see why "wealth heterogeneity" is not enough by itself, consider a CARA economy in which all investors share the same absolute risk aversion coefficient $\bar{\rho}$, but differ in initial wealth W_{0j}. With CARA preferences (and no borrowing constraints), the optimal number of shares at $t = 1$ is

$$q_{1j} = \frac{1}{\bar{\rho}} \frac{\mathbb{E}_j[z] - rp}{\mathbb{V}_j[z]},$$

which does *not* depend on W_{0j}. Intuitively, CARA investors scale utility so that "one more share" is equally risky whether you are rich or poor, so wealth does not change

how many shares you want to hold.

Because wealth does not change optimal position size, it also does not change (i) the gain from entering the stock market or (ii) the marginal benefit of acquiring more precision. Therefore, participation and learning do not sort by wealth in this CARA benchmark: absent additional constraints, either everyone enters and/or learns, or no one does. In that sense, *heterogeneity in wealth without heterogeneity in effective risk tolerance does not deliver the paper's results.*

A robustness experiment: common wealth, heterogeneous risk aversion. Conversely, heterogeneity in risk aversion *without* heterogeneity in initial wealth *does* deliver the same mechanism.

Consider an economy identical to the baseline except that everyone starts with the same wealth, $W_{0j} = W_0$, and investors differ only in risk aversion ρ_j. Interpreting ρ_j as a (relative) CRRA coefficient and applying the same local CARA approximation at wealth W_0 yields

$$\rho_j(W_0) = \frac{\rho_j}{W_0} \qquad \Rightarrow \qquad T_j \equiv \rho_j(W_0)^{-1} = \frac{W_0}{\rho_j}.$$

So even though wealth is identical across investors, differences in ρ_j generate differences in risk tolerance T_j. Lower ρ_j (more risk tolerance) implies a larger desired risky position, a higher private value of information, and a stronger incentive to pay the fixed costs of entering and learning.

Appendix F shows that equilibrium can still be summarized by the same two aggregate objects as in the baseline: (i) the total risk-bearing capacity among participants,

$$t \equiv \int_{j:\, 1_{\text{part},j}=1} T_j \, dG,$$

and (ii) the total *informed* risk-bearing capacity (risk tolerance weighted by precision),

$$s \equiv \int_{j:\, h_j>0} T_j h_j \, dG,$$

where G is now the cross-sectional distribution of ρ_j (equivalently of T_j). The marginal value of precision is the same object $D(s,t)$ as in Proposition 1, and the extensive margins are still characterized by simple cutoffs—but now the cutoffs are in *risk tolerance* rather than wealth:

$$T^{\text{part}} = \frac{2F}{D(s,t)}, \qquad T^{\text{info}} = \frac{2c_1}{D(s,t)}, \qquad \frac{T^{\text{info}}}{T^{\text{part}}} = \frac{c_1}{F}.$$

The economic interpretation is identical to the baseline. When $c_1 > F$, investors split into three groups: (i) high ρ (low T): do not participate; (ii) intermediate T: participate but remain uninformed; (iii) low ρ (high T): participate and become informed, choosing higher precision when their information cost is lower.

Because the sorting variable is still "how much risk one wants to take," the comparative statics and inequality incidence ranking carry over one-for-one:

- Lower participation costs ($F \downarrow$) pull in the most risk-averse investors at the left margin and compress between-group dispersion.
- Lower information adoption costs ($c_1 \downarrow$) expand the informed group (a thicker right tail among participants) and raise between-group dispersion.
- Lower information intensity costs ($c_0 \downarrow$) increase precision primarily among those already informed, raising within-informed dispersion.

In short, the primitive requirement for the results is heterogeneity that scales optimal position size. Wealth under DARA/CRRA is one empirically natural source, but heterogeneous risk aversion, background risk, horizons, or information-processing ability can play the same role.

5.3 Extension 2: Delegation (pooling into funds)

This section shows that the main inequality ranking is robust when investors can delegate information acquisition to a competitive active manager and benefit from pooling.

I implement this idea using a parsimonious delegated asset-management setup in the spirit of Gârleanu and Pedersen (2018).[6] The goal is not to build a rich institutional model of intermediation. The goal is to isolate how technology-induced declines in (i) participation/search frictions and (ii) information costs reallocate who enters, who becomes informed, and who benefits.

All institutional details and derivations that map my setup into the competitive delegation framework are in Appendix G; the main text focuses on the key incidence margins and the resulting inequality ranking.

Choice set and the delegation window. Each investor chooses among four options: (i) stay out; (ii) enter and trade uninformed; (iii) enter and acquire precision individually (as in the benchmark); or (iv) enter and pay a fixed *search cost* $F_m \geq 0$ to find a manager and join a pooled fund. Managers choose a common signal precision for the pool and charge fees that (by competition) cover information costs.

Pooling creates a simple "club-good" benefit: by sharing the fixed components of information production, delegated learning becomes attractive for investors who would not find it worthwhile to produce information on their own. At the same time, very wealthy (or very risk-tolerant) investors eventually prefer to self-learn because the private benefit of tailoring precision to a large position outweighs the advantages of pooling. This generates a *delegation window*: a range of investors for whom delegation dominates trading uninformed, but not necessarily self-learning.

What changes relative to the benchmark. Delegation changes *how* information is purchased, but it does not change the model's basic price-formation logic. The delegated sector adds to aggregate private precision s, so a larger delegated sector makes prices more informative. Compared to the benchmark, there are now three

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

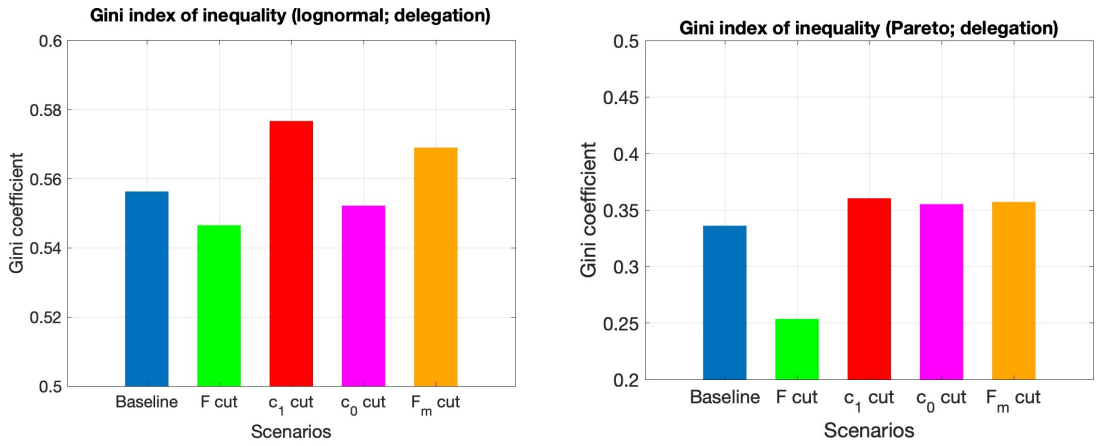
cutoffs: a participation cutoff W_0^{part}, a delegation cutoff W_0^{del} above which investors prefer to delegate rather than remain uninformed, and a switch cutoff W_0^{switch} above which investors prefer to self-learn rather than delegate.

How technology affects inequality with delegation. With delegation, technological improvements still map into inequality through the same three intuitive margins: (i) who enters at the participation cutoff, (ii) who moves from uninformed participation into (delegated) information, and (iii) who moves from delegation into self-learning at the top.

The overall effect on the Lorenz curve or the Gini depends on how much mass moves at each cutoff (which is determined by the wealth density) and on the general-equilibrium feedback through (s, t), which shifts the marginal value of information $D(s, t)$. Because both the masses and the feedback are endogenous, I do not claim a universal ranking between delegation and self-learning in all parameterizations. The welfare-distribution decomposition used in the quantitative exercises is summarized in Appendix G.5.2.

The qualitative ranking mirrors the benchmark. Lower participation costs ($F \downarrow$) expand entry at the left margin and typically compress inequality (a between-group compression effect). Lower information adoption costs ($c_1 \downarrow$) increase the mass of informed investors (delegated and self), thicken the right tail among participants, and raise inequality among participants. Lower information intensity costs ($c_0 \downarrow$) mainly scales up precision for agents who are already informed (delegated and self), so its aggregate inequality effects are typically smaller and calibration-dependent. Lower manager search costs ($F_m \downarrow$) expand delegation by making it cheaper to access pooled information. This pushes some uninformed participants into delegation, but it can also reduce entry indirectly by changing equilibrium informativeness and thus the value of entering. Which force dominates is theoretically ambiguous and is therefore quantified

Figure 3: **Inequality across scenarios (delegation extension).** Left: Lognormal prior. Right: Pareto prior. 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). 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.



Notes. Baseline (blue): lognormal (left) and Pareto (right) 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 with lognormal prior: 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). Reported Ginis with Pareto prior: 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.

in the calibration.

Figure 3 reports how inequality responds to each cost reduction in the delegation extension under the lognormal prior (left) and Pareto prior (right). Across priors, the *ranking* above is robust. The magnitude of delegation-driven changes is modest relative to the benchmark, and the main movements occur in the right tail when information frictions fall. Full derivations, cutoff definitions, and additional tables and figures are in Appendix G (Figures 13-14; Tables 8-9).

6 Conclusion

This paper asks a simple question: when financial information technology improves, who actually benefits? The answer is that the distributional effects depend on *which* friction technology relaxes. In the model, "FinTech" is not a single shock. It can make it cheaper for households to *enter* markets, or it can make it cheaper to *produce and exploit information*. Those changes have very different implications for inequality.

When technology reduces participation frictions—account opening, minimum balances, routine transaction and administrative costs—more households enter the stock market. Risk is shared more broadly, the gains from diversification reach further down the distribution, and wealth inequality falls *provided that* entry occurs into products and trading environments with basic safeguards (e.g., transparent fees and execution, simple menus, and limited conflicts of interest). In that sense, lowering participation costs can be equalizing, but only if the "front door" into markets is not a funnel into complexity and rent extraction.

When technology reduces information frictions—cheaper data, cheaper signals, cheaper processing and analytics—it raises the private returns to being informed. This reallocates surplus toward investors who are already positioned to acquire and use information, and it can widen wealth gaps. In equilibrium, cheaper information can also make uninformed participation less attractive by worsening the risk–return tradeoff faced by those who enter but do not (or cannot) become informed.

The quantitative illustration makes the contrast stark. In a high-participation baseline, a large reduction in participation costs lowers the wealth Gini by about 12%, whereas an equally large reduction in information costs raises it by about 22%. These incidence patterns are robust across alternative wealth distributions and across alternative targets for the shares of participants and informed investors.

These results also clarify why "more participation" is not automatically a cure for

inequality. If new entrants face opaque fees, complex products, or conflicted intermediation, part of the gains from entry can be diverted away from households and toward sophisticated counterparties or intermediaries—raising dispersion in realized outcomes rather than reducing it. The practical lesson is aligned with the broader "promise–peril" view of household finance: technologies that expand access can also be used to steer, extract rents, or amplify informational advantages.

Policy follows from this decomposition. Efforts to broaden participation—simpler onboarding, lower minimum balances, and defaults into low-fee diversified options—should be paired with protections that limit rent extraction: fee transparency, conflict management, suitability and advertising standards, and plain-language disclosure. By contrast, subsidizing data and advanced signals, or promoting sophisticated strategies for retail investors without comparable safeguards, risks entrenching top-end advantages by increasing the payoff to sophistication.

Finally, the paper is deliberately about *distributional incidence* rather than welfare. I use noisy supply in reduced form to generate trade and a role for information, and I do not model the welfare of noise traders (or other sources of non-fundamental trade) in a way that would support a full cost–benefit evaluation of "more information" versus "more participation." A natural next step is to embed the participation–information decomposition into a richer environment with endogenous non-fundamental trading, intermediary market power, and household mistakes, so that one can jointly assess efficiency, welfare, and inequality when new financial technologies arrive.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process. During the preparation of this work the author used ChatGPT4 in order to improve readability and grammar. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

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.

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

Bachas, Pierre, Paul Gertler, Sean Higgins, and Enrique Seira, "How Debit Cards Enable the Poor to Save More," *The Journal of Finance*, August 2021, *76* (4), 1913–1957.

Banerjee, Snehal, Jesse Davis, and Naveen Gondhi, "When Transparency Improves, Must Prices Reflect Fundamentals Better?," *The Review of Financial Studies*, 2018, *31* (6), 2377–2414.

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.

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.

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

Bretscher, Lorenzo, Francisco J. Gomes, Riccardo Sabbatucci, and Andrea Tamoni, "Return Heterogeneity in Retirement Accounts," Technical Report, SSRN, Stockholm, Sweden 2022.

Calvet, Laurent E, Laurent Bach, and Paolo Sodini, "Rich Pickings? Risk, Return, and Skill in the Portfolios of the Wealthy," HEC Research Papers Series 1126, HEC Paris January 2019.

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.

_ , 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.

Campbell, John Y. and Tarun Ramadorai, *Fixed: Why Personal Finance Is Broken and How to Make It Work for Everyone*, Princeton University Press, 2025.

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, 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.

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

Fedyk, Anastassia, Roxana Mihet, Kumar Rishabh, and Orlando Gomes, "Data Innovation Complementarity and Firm Growth," Technical Report 23-86, SFI September 2023.

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

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.

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

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).

Gomes, Francisco J. and Alexander Michaelides, "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence," *The Journal of Finance*, April 2005, *60* (2), 869–904.

_ , Michael Haliassos, and Tarun Ramadorai, "Household Finance," *Journal of Economic Literature*, 2021, *59* (3), 919–1000.

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.

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 Painter, Panos 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.

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

Martineau, Charles and Marius Zoican, "Retail Trading and Analyst Coverage," *Journal of Financial Markets*, 2023, *66*, 100849.

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, Kumar Rishabh, and Orlando Gomes, "Is it AI or data that drives firm market power?," *Journal of Monetary Economics*, 2026, *157*, 103878.

Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar, "Building State Capacity: Evidence from Biometric Smartcards in India," *American Economic Review*, October 2016, *106* (10), 2895–2929.

Nieuwerburgh, Stijn Van and Laura Veldkamp, "Information Immobility and the Home Bias Puzzle," *The Journal of Finance*, 2009, *64* (3), 1187–1215.

__ and __, "Information Acquisition and Under-Diversification," *The Review of Economic Studies*, 2010, *77* (2), 779–805.

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.

Philippon, Thomas, "On Fintech and Financial Inclusion," BIS Working Papers 841, Bank for International Settlements February 2020. Manuscript dated September 2019.

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 Under Heterogeneous Returns," *The Quarterly Journal of Economics*, 02 2023, *138* (1), 515–573.

Suri, Tavneet and William Jack, "The long-run poverty and gender impacts of mobile money," *Science*, December 2016, *354* (6317), 1288–1292.

Veldkamp, Laura, "Information Markets and the Comovement of Asset Prices," *Review of Economic Studies*, 2006, *73* (3), 823–845.

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

Vissing-Jørgensen, Annette, "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy*, August 2002, *110* (4), 825–853.

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 1: 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; $\text{Gini} = \dfrac{1}{2\mu_2}\,\mathbb{E}[\,\lvert W_2 - W_2'\rvert\,]$
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

(continued on next page)

(continued from previous page)

Symbol	Definition / interpretation
$\mathcal{I}_j$	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
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

(continued on next page)

(continued from previous page)

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

B Baseline Model Derivations

B.1 Portfolio Demand

Proposition 2. *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

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

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 3. *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[7]

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

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[8]

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

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

[7] 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.

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

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

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

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

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[9]

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

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

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

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

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

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

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.

Price precision. Plugging in the coefficients from equations B.14a, B.14b 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*

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

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

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

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 2. *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,

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

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

Simplify the numerator.

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

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 3. *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 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 Appendix B.4, $D = \dfrac{\sigma_x^2 + 2st + t^2 h_0}{t^2(h_0 + s^2/\sigma_x^2)^2}$.

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.6 Existence and local uniqueness of (s,t)

Proposition 4 (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.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 *B*.10.

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

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:

$$W_0 < W_0^{\text{part}} \implies \text{non-participant};$$

$$W_0^{\text{part}} \le W_0 < W_0^{\text{info}} \implies \text{uninformed participant};$$

$$W_0 \ge W_0^{\text{info}} \implies \text{informed participant}.$$

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

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

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

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

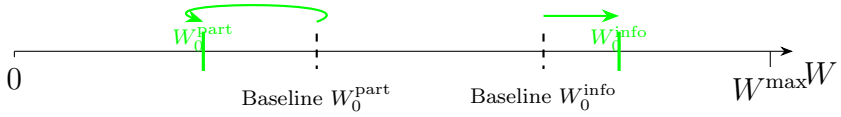
B.9.1 Lemma 1.1. Entry fee

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

Proof. Recall the cutoff formulas (**??**):

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

Figure 4: **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).



Let

$$D_F \;\equiv\; \frac{d}{dF} D\big(s(F), t(F)\big) \;=\; D_s(s,t)\, s_F \;+\; D_t(s,t)\, t_F.$$

Using (B.32), the total derivatives satisfy

$$\frac{dW_0^{\text{part}}}{dF} = \frac{2\rho}{D} - \frac{2\rho F}{D^2} D_F = \frac{2\rho}{D}\Big(1 - \frac{F}{D} D_F\Big),$$
$$\frac{dW_0^{\text{info}}}{dF} = -\frac{2\rho c_1}{D^2} D_F.$$

Sufficient condition. Assume that, in the active-information region,

$$0 \;<\; D_F \;<\; \frac{D}{F} \qquad \text{equivalently} \qquad 0 \;<\; \frac{d\ln D}{d\ln F} \;<\; 1. \tag{B.33}$$

Then $D_F > 0$ implies $\frac{dW_0^{\text{info}}}{dF} < 0$, so a *decrease* in F raises W_0^{info}. Moreover, (B.33) implies $1 - \frac{F}{D} D_F > 0$, hence $\frac{dW_0^{\text{part}}}{dF} > 0$, so a *decrease* in F lowers W_0^{part}. This yields $W_0^{\text{part}} \downarrow$ and $W_0^{\text{info}} \uparrow$ when $F \downarrow$. □

B.9.2 Lemma 2. Linear c_1 information cost

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

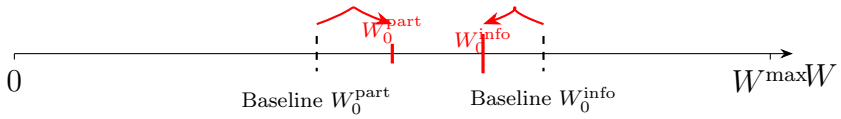


Figure 5: **Threshold movements for c_1 cut.** Baseline thresholds W_0^{part} and W_0^{info} (dashed gray) on the wealth axis $[0, W^{\max}]$. When c_1 falls, the thresholds move *closer together* (red).

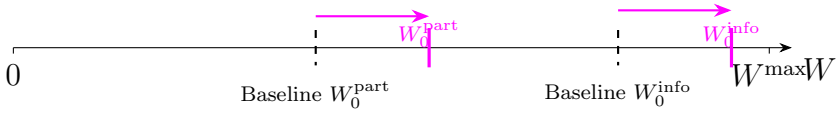


Figure 6: **Threshold movements for c_0 cut.** Baseline thresholds W_0^{part} and W_0^{info} (dashed gray) on the wealth axis $[0, W^{\max}]$. When c_0 falls, the thresholds shift *together to the right* (magenta).

Proof. Recall (B.7):

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

Let

$$D_{c_1} \;\equiv\; \frac{d}{dc_1} D\big(s(c_1), t(c_1)\big) \;=\; D_s(s,t)\, s_{c_1} \;+\; D_t(s,t)\, t_{c_1}.$$

Using (B.32), the total derivatives satisfy

$$\frac{dW_0^{\text{info}}}{dc_1} = \frac{2\rho}{D} - \frac{2\rho c_1}{D^2} D_{c_1} = \frac{2\rho}{D}\Big(1 - \frac{c_1}{D} D_{c_1}\Big),$$
$$\frac{dW_0^{\text{part}}}{dc_1} = -\frac{2\rho F}{D^2} D_{c_1}.$$

Sufficient condition. Assume that, in the active-information region,

$$0 \;<\; D_{c_1} \;<\; \frac{D}{c_1} \qquad \text{equivalently} \qquad 0 \;<\; \frac{d\ln D}{d\ln c_1} \;<\; 1. \tag{B.34}$$

Then $D_{c_1} > 0$ implies $\frac{dW_0^{\text{part}}}{dc_1} < 0$, so a *decrease* in c_1 raises W_0^{part}. Moreover, (B.34) implies $1 - \frac{c_1}{D} D_{c_1} > 0$, hence $\frac{dW_0^{\text{info}}}{dc_1} > 0$, so a *decrease* in c_1 lowers W_0^{info}. This yields $W_0^{\text{info}} \downarrow$ and $W_0^{\text{part}} \uparrow$ when $c_1 \downarrow$. □

B.9.3 Lemma 3. Quadratic c_0 information cost

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

Proof (direction-free). Holding (F, c_1) fixed, the cutoffs satisfy (B.7)

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

Let $D_{c_0} \equiv \frac{d}{dc_0} D(s(c_0), t(c_0))$. By the chain rule,

$$D_{c_0} \;=\; D_s(s,t)\, s_{c_0} \;+\; D_t(s,t)\, t_{c_0}.$$

Differentiating the cutoffs yields

$$\frac{dW_0^{\text{part}}}{dc_0} = -\frac{2\rho F}{D(s,t)^2}\, D_{c_0}, \qquad \frac{dW_0^{\text{info}}}{dc_0} = -\frac{2\rho c_1}{D(s,t)^2}\, D_{c_0}.$$

Hence W_0^{part} and W_0^{info} move in the *same* direction for any perturbation of c_0, and their ratio satisfies

$$\frac{W_0^{\text{info}}}{W_0^{\text{part}}} = \frac{c_1}{F}$$

independently of $D(s,t)$. This proves the direction-free statement. □

Proof (direction-conditional: a sufficient sign for D_{c_0}). Maintain the active-information region restrictions $D_s(s,t) < 0$ and $D_t(s,t) < 0$. Assume further:

(i) Intensive-margin effect on price informativeness. An increase in the intensity-cost parameter reduces equilibrium price informativeness:

$$s_{c_0} < 0. \tag{B.35}$$

(ii) Participation aggregates depend on c_0 only through the cutoffs. For fixed (F, c_1), the risk-bearing sector size can be written as

$$t \;=\; \mathcal{T}\big(W_0^{\text{part}},\, W_0^{\text{info}}\big)\,,$$

with $\mathcal{T}_{W_0^{\text{part}}} < 0$ and $\mathcal{T}_{W_0^{\text{info}}} < 0$ (i.e., raising either cutoff weakly reduces participation/risk-bearing).

Under (ii) and (**??**), t is a function of D alone when (F, c_1) are fixed:

$$t \;=\; \widetilde{\mathcal{T}}(D) \equiv \mathcal{T}\left(\frac{2\rho F}{D},\, \frac{2\rho c_1}{D}\right).$$

Differentiating gives

$$\widetilde{\mathcal{T}}'(D) = \mathcal{T}_{W_0^{\text{part}}}\Big(-\frac{2\rho F}{D^2}\Big) + \mathcal{T}_{W_0^{\text{info}}}\Big(-\frac{2\rho c_1}{D^2}\Big) \;>\; 0,$$

because $\mathcal{T}_{W_0^{\text{part}}} < 0$, $\mathcal{T}_{W_0^{\text{info}}} < 0$, and the cutoff derivatives in parentheses are negative. Therefore,

$$t_{c_0} = \widetilde{\mathcal{T}}'(D)\, D_{c_0}. \tag{B.36}$$

Substitute (B.36) into the definition of D_{c_0}:

$$D_{c_0} = D_s\, s_{c_0} + D_t\, t_{c_0} = D_s\, s_{c_0} + D_t\, \widetilde{\mathcal{T}}'(D)\, D_{c_0}.$$

Rearranging,

$$\Big(1 - D_t\, \widetilde{\mathcal{T}}'(D)\Big) D_{c_0} = D_s\, s_{c_0}.$$

Since $D_t < 0$ and $\widetilde{\mathcal{T}}'(D) > 0$, the prefactor satisfies $1 - D_t\, \widetilde{\mathcal{T}}'(D) > 0$. Hence

$$\operatorname{sign}(D_{c_0}) = \operatorname{sign}(D_s\, s_{c_0}).$$

Using $D_s < 0$ and $s_{c_0} < 0$ from (B.35), we obtain $D_{c_0} > 0$. Consequently, a *decrease* in c_0 lowers D and therefore raises both W_0^{part} and W_0^{info} (by (**??**)). □

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.

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

Consider an *uninformed* participant ($\sigma_{s,j}^{-2} = 0$). 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)

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

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

$$S_U^2(s,t) = \frac{\sigma_x^2 + 2st + t^2 h_0}{t^2\big(h_0 + s^2/\sigma_x^2\big)^2}. \tag{B.37}$$

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 (B.37). Therefore

$$D(s,t) \;=\; S_U^2(s,t) \quad \text{for all } s > 0,\ t > 0. \tag{B.38}$$

This concludes the proof. This proves that the marginal benefit of producing a bit of private information equals the expected squared Sharpe ratio available to an uninformed participant.

C Lorenz/Gini proofs for Theorem 1

This appendix proves Theorem 1. Throughout, fix an interior equilibrium and write $\theta = (F, c_1, c_0)$.

C.0.1 Mean-preserving normalization

Let $Y(W_0;\theta)$ denote certainty-equivalent terminal wealth under technology vector $\theta = (F, c_1, c_0)$ and initial-wealth distribution V. Let

$$\mu(\theta) \equiv \int Y(W_0;\theta)\, dV(W_0).$$

Fix a baseline θ_0 and define the scale-normalized mapping

$$\widetilde{Y}(W_0;\theta) \equiv \frac{\mu(\theta_0)}{\mu(\theta)}\, Y(W_0;\theta).$$

Then $\int \widetilde{Y}(W_0;\theta)\,dV(W_0) = \mu(\theta_0)$ for all θ. Because Lorenz curves and the Gini coefficient are scale-invariant, replacing Y by $\widetilde{Y}$ leaves $L(\cdot;\theta)$ and $G(\theta)$ unchanged, while ensuring $d\mu = 0$ in the Gateaux variation used in the proofs. In the quantitative illustration, I implement the same normalization by rescaling market payoffs so that $\mathbb{E}[\widetilde{Y}(W_0;\theta)]$ is constant across counterfactuals.

C.1 Preliminaries: Lorenz Gateaux variation and Dalton principle

Let $\widetilde{Y}(W_0;\theta)$ be the scale-normalized CE wealth defined in Appendix C.0.1, so that $\mu \equiv \int \widetilde{Y}(W_0;\theta)\,dV(W_0)$ is constant across counterfactuals.

Let $L(p;\theta)$ be the Lorenz curve of $\widetilde{Y}(\cdot;\theta)$ and $q(p)$ the p-quantile of W_0 under V. For a small perturbation $d\theta$, the Gateaux variation of the Lorenz curve is

$$dL(p;\theta) = \frac{1}{\mu}\int_0^{q(p)} d\widetilde{Y}(W_0;\theta)\,dV(W_0), \tag{C.1}$$

since $d\mu = 0$ by construction. The Gini coefficient is $G(\theta) = 1 - 2\int_0^1 L(p;\theta)\,dp$, hence

$$dG(\theta) = -2\int_0^1 dL(p;\theta)\,dp. \tag{C.2}$$

By the Dalton principle: if $d\widetilde{Y}(\cdot;\theta)$ is *progressive* (it raises lower-quantile wealth more than upper-quantile wealth, so that the cumulative integral in (C.1) is nonnegative for all p), then $dL(p;\theta) \geq 0$ and $dG(\theta) \leq 0$. If it is *regressive*, then $dL(p;\theta) \leq 0$ and $dG(\theta) \geq 0$.

C.2 Shape of the CE schedule

Using (13) in the main text, for fixed equilibrium objects (hence fixed D),

$$\Delta U^{\text{mkt}}(W_0;\theta) = \left(\frac{W_0}{2\rho}D(\theta) - F\right)\mathbf{1}\{W_0 \geq W_0^{\text{part}}(\theta)\}, \tag{C.3}$$

$$\Delta U^{\text{info}}(W_0;\theta) = \frac{1}{2c_0}\left(\frac{W_0}{2\rho}D(\theta) - c_1\right)^2\mathbf{1}\{W_0 \geq W_0^{\text{info}}(\theta)\}. \tag{C.4}$$

Hence $\widetilde{Y}(W_0;\theta)$ is weakly increasing in W_0 everywhere; moreover, on the informed region $[W_0^{\text{info}}(\theta),\infty)$ it is strictly convex in W_0 because ΔU^{info} is quadratic in W_0.

C.3 Proof of Theorem 1(1): $F \downarrow$

Consider a small reduction $dF < 0$, holding (c_1, c_0) fixed. By Lemma 1 in the main text, the cutoff moves as $dW_0^{\text{part}} < 0$ and $dW_0^{\text{info}} > 0$. There are two first-order regime changes:

- *Entry expansion at the bottom:* types in $[W_0^{\text{part}}(\theta)+dW_0^{\text{part}},\ W_0^{\text{part}}(\theta))$ switch from outside to market participation. Their CE wealth rises discretely by the (positive) net surplus from entering at the margin. This component of $d\widetilde{Y}$ is supported at the bottom of the distribution and is therefore progressive.

- *Information retrenchment at the top:* types in $[W_0^{\text{info}}(\theta),\ W_0^{\text{info}}(\theta)+dW_0^{\text{info}})$ switch from informed to uninformed. They lose the information rent $\Delta U^{\text{info}}(W_0;\theta)$, which is strictly increasing in W_0 on the informed region. This component of $d\widetilde{Y}$ is supported in the upper tail and is therefore also progressive (it reduces upper-tail outcomes relative to the rest).

Any within-regime effect of F operates through a uniform reduction in the fixed entry payment for participants and is offset by the scalar $\kappa(\theta)$ under the mean-preserving normalization. Combining the two regime-change components implies that $d\widetilde{Y}(\cdot;\theta)$ is

progressive in the sense of (C.1). Therefore $dL(p;\theta) \geq 0$ for all p and $dG(\theta) \leq 0$ by (C.2). This proves part (1).

C.4 Proof of Theorem 1(2): $c_1 \downarrow$

Consider a small reduction $dc_1 < 0$, holding (F, c_0) fixed. By Lemma 2 in the main text, $dW_0^{\text{part}} > 0$ and $dW_0^{\text{info}} < 0$. There are again two first-order regime changes:

- *Exit at the bottom:* types in $[W_0^{\text{part}}(\theta),\ W_0^{\text{part}}(\theta) + dW_0^{\text{part}})$ switch from market participation to the outside option. They lose $\Delta U^{\text{mkt}}(W_0;\theta)$, which is positive and increasing in W_0 on the participants' support. This is a regressive change because it reduces outcomes in the lower tail.

- *Adoption at the top:* types in $[W_0^{\text{info}}(\theta)+dW_0^{\text{info}},\ W_0^{\text{info}}(\theta))$ switch from uninformed to informed. They gain $\Delta U^{\text{info}}(W_0;\theta)$, which is strictly increasing in W_0 on the informed region. This is a regressive change because it raises outcomes in the upper tail.

Moreover, even holding cutoffs fixed, lowering c_1 raises the information rent pointwise for informed investors:

$$\frac{\partial}{\partial c_1}\Delta U^{\text{info}}(W_0;\theta) = -\frac{1}{c_0}\left(\frac{W_0}{2\rho}D(\theta) - c_1\right)\mathbf{1}\{W_0 \geq W_0^{\text{info}}(\theta)\} < 0,$$

so $dc_1 < 0$ induces an *increasing* perturbation in W_0 on the informed region, which is regressive after mean normalization. Hence $d\widetilde{Y}(\cdot;\theta)$ is regressive overall, which implies $dL(p;\theta) \leq 0$ and $dG(\theta) \geq 0$. This proves part (2).

C.5 Proof of Theorem 1(3): $c_0 \downarrow$

Consider $dc_0 < 0$, holding (F, c_1) fixed. Conditional on being informed, the information rent satisfies

$$\frac{\partial}{\partial c_0}\Delta U^{\text{info}}(W_0;\theta) = -\frac{1}{2c_0^2}\left(\frac{W_0}{2\rho}D(\theta) - c_1\right)^2 \mathbf{1}\{W_0 \geq W_0^{\text{info}}(\theta)\} < 0,$$

so $dc_0 < 0$ increases $\Delta U^{\text{info}}(W_0;\theta)$ by an amount that is strictly increasing and strictly convex in W_0 on the informed region. Under mean normalization restricted to the informed subpopulation, this is a regressive change, so within-informed Lorenz shifts down and the conditional Gini increases.

For the *aggregate* Lorenz/Gini, c_0 also shifts equilibrium objects (s, t) and hence $D(s, t)$, changing both cutoffs in (14) and the linear component ΔU^{mkt}. This general-equilibrium channel can partially offset the regressive within-informed convexity effect, so the overall sign need not be pinned down without additional restrictions. This proves part (3).

D Quantitative Illustration: Calibration, Computation, and Additional Results

This appendix documents the numerical implementation underlying Section 4. The exercise is a *quantitative illustration* of the comparative-statics and incidence results, not an estimation: the technology parameters (F, c_1, c_0) are not directly observable and the model is intentionally stylized. The objective is therefore (i) transparency and replicability of the fixed-point computation, and (ii) internal-consistency checks that map directly into the theory (cutoff movements, group shares, and within-informed dispersion), rather than fit to a large set of moments.

Throughout, I work with the certainty-equivalent (CE) terminal-wealth mapping $Y(W_0; \theta)$ implied by the equilibrium in Section 2, and I report Lorenz/Gini objects for the cross-section of $Y(W_0; \theta)$ after a mean-preserving normalization (see Appendix C.0.1).

D.1 Wealth priors and truncated moments

I treat the initial wealth distribution $W_0 \sim V$ (with density g) as primitive and report results under three priors:

1. **Lognormal:** a convenient benchmark with substantial mass in the middle.
2. **Pareto (upper-tail):** a heavy-tail benchmark, motivated by macro-inequality evidence.
3. **Uniform (thin-tail):** a bounded-support benchmark reported as a theoretical check that isolates the role of tail thickness in the magnitude of inequality responses. Because the Uniform prior mechanically removes extreme tails, it shrinks the range of attainable Lorenz/Gini movements while preserving the *ranking* across technology levers.

For the fixed-point system below, it is convenient to work with truncated moments.

For any upper support bound $\bar{W} \in (0, \infty]$ (with $\bar{W} = \infty$ for Lognormal), define

$$M_k(u) \equiv \int_u^{\bar{W}} W^k \, g(W) \, dW, \qquad k \in \{1, 2\}.$$

Closed forms used in the computations are:

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

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

Truncated Pareto ($W \in [w_{\min}, W_{\max}]$). Let $\alpha > 1$ and define the truncated Pareto density on $[w_{\min}, W_{\max}]$,

$$g(W) = \frac{\alpha \, w_{\min}^{\alpha}}{1 - (w_{\min}/W_{\max})^{\alpha}} \, W^{-(\alpha+1)}.$$

For $u_+ \equiv \max\{u, w_{\min}\}$, the truncated moments are

$$M_1(u) = \frac{\alpha \, w_{\min}^{\alpha}}{(\alpha - 1)\left[1 - (w_{\min}/W_{\max})^{\alpha}\right]} \left(u_+^{1-\alpha} - W_{\max}^{1-\alpha}\right),$$

and

$$M_2(u) = \begin{cases} \dfrac{\alpha \, w_{\min}^{\alpha}}{(\alpha - 2)\left[1 - (w_{\min}/W_{\max})^{\alpha}\right]} \left(u_+^{2-\alpha} - W_{\max}^{2-\alpha}\right), & \alpha \neq 2, \\ \dfrac{\alpha \, w_{\min}^{2}}{1 - (w_{\min}/W_{\max})^{2}} \, \ln\Big(\dfrac{W_{\max}}{u_+}\Big), & \alpha = 2. \end{cases}$$

This truncation is both numerically convenient and economically interpretable: it corresponds to a finite upper tail in finite samples and prevents divergence of second moments when $\alpha \leq 2$.

Uniform ($W \sim \text{Unif}[0, \bar{W}]$). Here $g(W) = 1/\bar{W}$ and for $u \in [0, \bar{W}]$,

$$M_1(u) = \frac{\bar{W}^2 - u^2}{2\bar{W}}, \qquad M_2(u) = \frac{\bar{W}^3 - u^3}{3\bar{W}},$$

and $M_k(u) = 0$ for $u \geq \bar{W}$.

D.2 Equilibrium system in truncated-moment form

Fix technology parameters $\theta = (F, c_1, c_0)$ and primitives $(\rho, \sigma_x^2, \sigma_z^2, h_0, r)$ as in Section 2. With absolute risk tolerance $r(W) = W/\rho$ and quadratic-plus-linear information costs, the optimal precision for an informed investor with wealth $W \geq W_0^{\text{info}}$ is

$$\sigma_s^{-2}(W) = \frac{(W/\rho)\, D(s,t) - 2c_1}{2c_0}, \tag{D.1}$$

so that the information-adoption cutoff solves $\sigma_s^{-2}(W_0^{\text{info}}) = 0$.

Let t denote aggregate risk tolerance of participants and s denote aggregate "information-weighted" risk tolerance of informed traders (as defined in the main text). The equilibrium can be written as the fixed point in (s,t):

$$\Psi_1(s,t) \equiv t - \frac{1}{\rho}\, M_1\big(W_0^{\text{part}}\big) = 0, \tag{D.2}$$

$$\Psi_2(s,t) \equiv s - \frac{D(s,t)}{2\rho^2 c_0}\, M_2\big(W_0^{\text{info}}\big) + \frac{c_1}{c_0 \rho}\, M_1\big(W_0^{\text{info}}\big) = 0, \tag{D.3}$$

where the wealth cutoffs are

$$W_0^{\text{part}} = \frac{2\rho F}{D(s,t)}, \qquad W_0^{\text{info}} = \frac{2\rho c_1}{D(s,t)}, \qquad \frac{W_0^{\text{info}}}{W_0^{\text{part}}} = \frac{c_1}{F}, \tag{D.4}$$

and the (uninformed) squared Sharpe object is

$$D(s,t) = \frac{\sigma_x^2 + 2st + t^2 h_0}{t^2\big(h_0 + s^2/\sigma_x^2\big)^2}. \tag{D.5}$$

Equations (D.2)–(D.3) are solved by Newton's method with backtracking (details and Jacobian in Appendix D.3).

D.3 Newton solver and Jacobian

Let $\vartheta = (s,t)'$ and $\Psi(\vartheta) = (\Psi_2(\vartheta), \Psi_1(\vartheta))'$. The Newton step is $\vartheta \leftarrow \vartheta - \alpha\, J(\vartheta)^{-1}\Psi(\vartheta)$ with backtracking $\alpha \in (0,1]$ chosen to decrease $\|\Psi\|$.

The Jacobian is

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

Write $u_{\mathrm{p}} = W_0^{\mathrm{part}}$ and $u_{\mathrm{i}} = W_0^{\mathrm{info}}$. Using Leibniz' rule and the threshold property $\sigma_s^{-2}(u_{\mathrm{i}}) = 0$, one obtains

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

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

with

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

and derivatives of D computed from (D.5):

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

Stopping criterion. I stop when $\|\Psi(\vartheta)\|_\infty < 10^{-10}$, and I verify that the Jacobian is nonsingular at the solution across all reported calibrations.

D.4 From initial wealth to CE terminal wealth

Given the fixed point (s, t) and cutoffs (D.4), each investor is assigned to a state as a function of W_0:

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

For each state, I compute the certainty-equivalent terminal wealth $Y(W_0; \theta)$ using the value expressions in Section 2 (i.e., expected trading gains net of participation and information costs under mean-variance preferences). This produces a *deterministic*

mapping $W_0 \mapsto Y(W_0;\theta)$, which is the object used in the Lorenz/Gini computations.

Mean-preserving normalization. To isolate incidence, I apply a common multiplicative rescaling to market payoffs so that $\mathbb{E}[Y(W_0;\theta)]$ is constant across scenarios; see Appendix C.0.1.

D.5 Inequality objects and numerical computation

Given $Y(W_0;\theta)$ and a draw $\{W_0^{(i)}\}_{i=1}^N$ from V, form $Y^{(i)} = Y(W_0^{(i)};\theta)$ and compute:

Lorenz curve. Sort $\{Y^{(i)}\}$ ascending to $\{Y^{(i:N)}\}$ and define cumulative shares

$$L(p) = \frac{\sum_{i=1}^{\lfloor pN \rfloor} Y^{(i:N)}}{\sum_{i=1}^{N} Y^{(i:N)}}, \qquad p \in [0,1].$$

Gini coefficient. Using the standard sample formula,

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

I also compute conditional (within-informed) Lorenz/Gini objects by restricting to the subpopulation $\{i : W_0^{(i)} \geq W_0^{\text{info}}(\theta)\}$, consistent with the decomposition in the main text.

D.6 Baseline targets and counterfactual design

Targets and normalizations. For each wealth prior V, I discipline the *extensive margins* emphasized by the theory by matching (i) a baseline participation share $\Pr(W_0 \geq W_0^{\text{part}}) = 0.60$ and (ii) a baseline informed share $\Pr(W_0 \geq W_0^{\text{info}}) = 0.05$. These targets pin down the *locations* of the endogenous cutoffs as quantiles of V, and thus discipline the two fixed-cost components (F, c_1).

The remaining technology parameter c_0 governs the *intensity* (curvature) of information production and primarily affects *within-informed dispersion*. In the quantitative

illustration, I treat c_0 as an *intensity-scale normalization*: when feasible, I choose c_0 so that the baseline model-implied Gini of certainty-equivalent terminal wealth matches a reference level; when the reference level lies outside the attainable set given (i)–(ii), I select the c_0 that minimizes the distance to the reference and report the achieved baseline Gini.[10] This step is not interpreted as estimation; it simply anchors the magnitude of the intensive-margin wedge while leaving the incidence *rankings* (e.g., Gini coefficients and Lorenz-curves) unchanged.

Counterfactual technology shocks. Starting from the baseline, I implement one-at-a-time reductions in:

$$F \downarrow \quad \text{(entry/onboarding)}, \qquad c_1 \downarrow \quad \text{(information adoption)}, \qquad c_0 \downarrow \quad \text{(information intensity)}.$$

The benchmark counterfactuals in the main text are 70% reductions in each cost; additional magnitudes are reported in Appendix D.7.

Internal consistency checks. For each counterfactual, I compute: (i) the direction of cutoff movements (as predicted by Lemma 1), (ii) the induced group shares (π_U, π_I), and (iii) Lorenz/Gini and intensity proxies. These diagnostics are the main reason to include numerical material in a theory paper: they verify that the simulated distributional movements align with the analytic incidence decomposition.

Computation workflow (replicable steps). For each scenario θ (baseline and each one-at-a-time technology shock), the numerical implementation follows the same steps:

1. **Solve equilibrium.** Solve the fixed-point system (D.2)–(D.3) for $(s(\theta), t(\theta))$ using Newton with backtracking as described in Appendix D.3. Compute $D(s,t)$ and the

[10] Given the extensive-margin targets and the mean-preserving normalization, varying c_0 can shift aggregate Lorenz/Gini objects only through within-informed dispersion and general-equilibrium feedback in $D(s,t)$; as a result, the attainable range of baseline Ginis can be narrow, especially under thin-tailed priors.

implied cutoffs $(W_0^{\text{part}}(\theta), W_0^{\text{info}}(\theta))$ from (D.4).

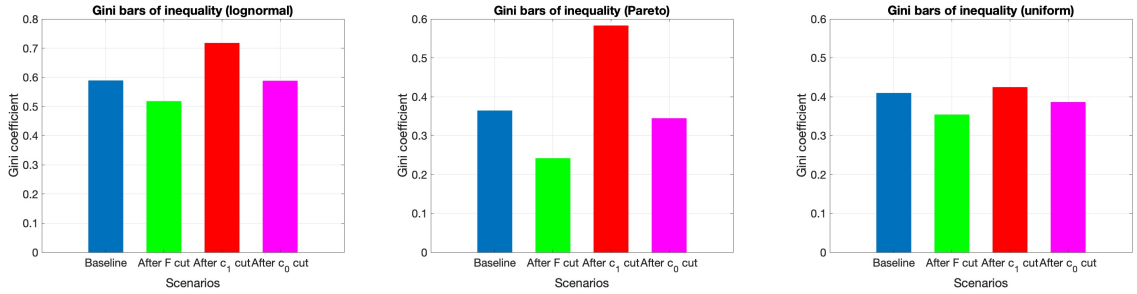
2. **Construct the CE mapping.** Using the equilibrium objects and cutoffs, evaluate the certainty-equivalent terminal-wealth mapping $Y(W_0;\theta)$ as described in Appendix D.4.
3. **Apply mean-preserving normalization.** Rescale the mapping to $\widetilde{Y}(W_0;\theta)$ using Appendix C.0.1 so that the cross-scenario mean is constant and distributional comparisons isolate incidence.
4. **Compute inequality objects.** Using either a large Monte Carlo draw from V or a fine deterministic grid over W_0, compute Lorenz curves and Gini coefficients for $\widetilde{Y}(\cdot;\theta)$ as in Appendix D.5. Compute the conditional (within-informed) Lorenz/Gini by restricting to $\{W_0 \geq W_0^{\text{info}}(\theta)\}$.
5. **Record decomposition diagnostics.** Record group shares (π_U, π_I) and within-informed dispersion/intensity diagnostics used in Section 4.3. As a numerical check, verify that cutoff movements across scenarios match Lemma 1.

D.7 Supplementary figures and robustness tables

Appendix Figures 7–9 report Lorenz/Gini objects and terminal-wealth densities under Lognormal, Pareto, and Uniform priors using a piecewise linear approach computation.

Figure 7: **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).

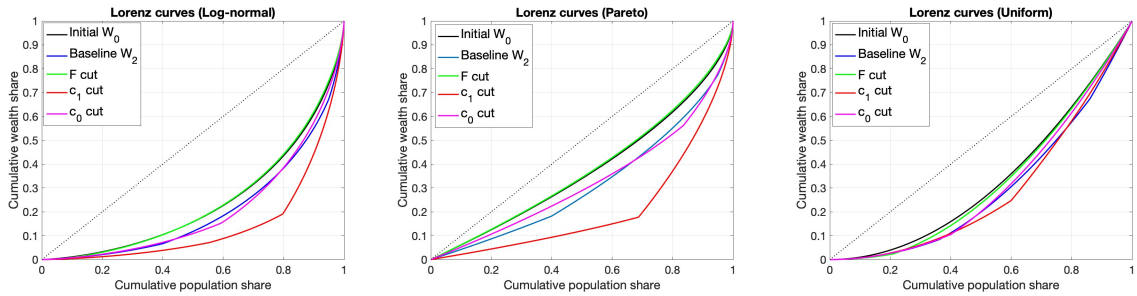
Left. Log-normal Middle. Pareto Right. Uniform



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.

Figure 8: **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).

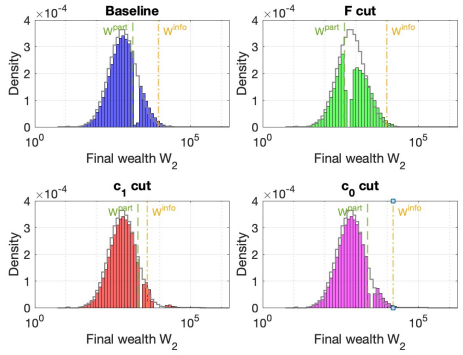
Left. Log-normal Middle. Pareto Right. Uniform



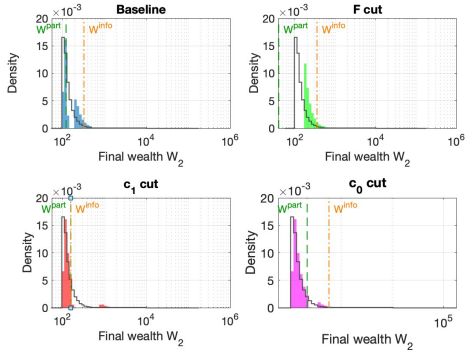
Notes. Initial wealth W_0 and primitives are held fixed across scenarios; in each case we solve the fixed point for (s, t). Patterns are consistent with the comparative statics independently of the distribution used: the F cut compresses inequality, the c_1 cut raises it, and the c_0 cut leaves it broadly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

Figure 9: **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.

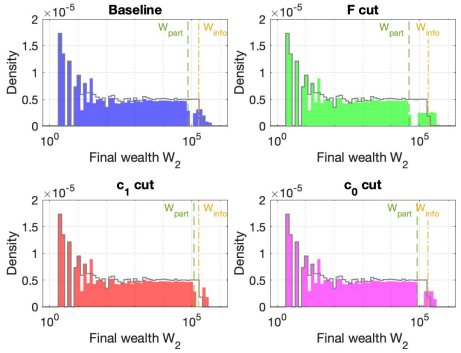
Left. Log-normal



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.

D.7.1 Alternative shares of participation and information.

Appendix Tables 2–4 report robustness to alternative baseline participation/informed shares and alternative magnitudes of cost cuts.

Table 2: 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.

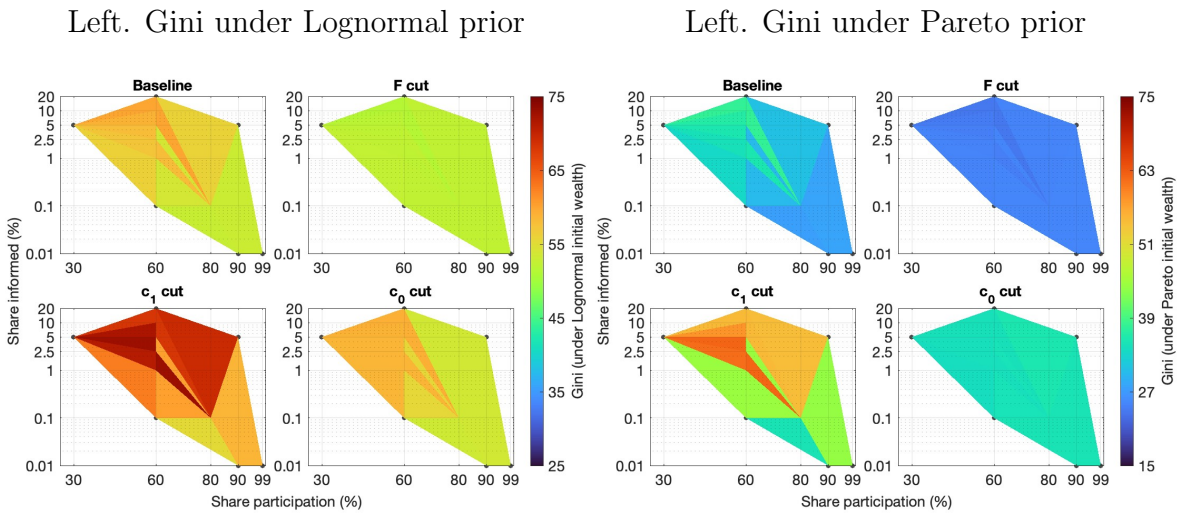
Table 3: Gini Coefficient (Pareto): Robustness for different levels of shares of participating and informed investors.

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 2 (for lognormal) and 3 (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 10 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 10: **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.

Figure 10 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 10 highlights the role of tail 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.

D.7.2 Alternative magnitudes of technology improvement shocks

Table 4: 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.

Table 4 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 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.

D.7.3 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 5 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

Table 5: **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 ($\times 10^{-8}$)	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
c1_cut	119.56	91.96	1.32	159.46	159.46	31.40%	31.40%	0.583
c0_cut	147.70	62.22	1.02	206.89	559.01	16.20%	1.40%	0.345
	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.

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.

E Relation to Empirical 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. My 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, Aymard and Cenesizoglu, 2025), robo-advisors expanded access to diversified portfolios (D'Acunto, Prabhala and Rossi, 2019), and digital financial inclusion lifted the lower tail, reducing extreme poverty, raising net benefits to poor households, and increasing formal savings (Suri and Jack, 2016; Muralidharan, Niehaus and Sukhtankar, 2016; Bachas, Gertler, Higgins and Seira, 2021). Descriptive evidence from the Survey of Consumer Finances (SCF) confirms that participation remains sharply skewed by wealth. Table 6 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 6: 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, Gambacorta and Mihet (2023) also supports this channel: improvements in FinTech infrastructure change portfolio allocations and participation patterns in ways consistent with the comparative statics.

Long-term Wealth Dispersion Trends These mechanisms are consistent with broader wealth concentration dynamics. Table 7 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.

Table 7: **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.

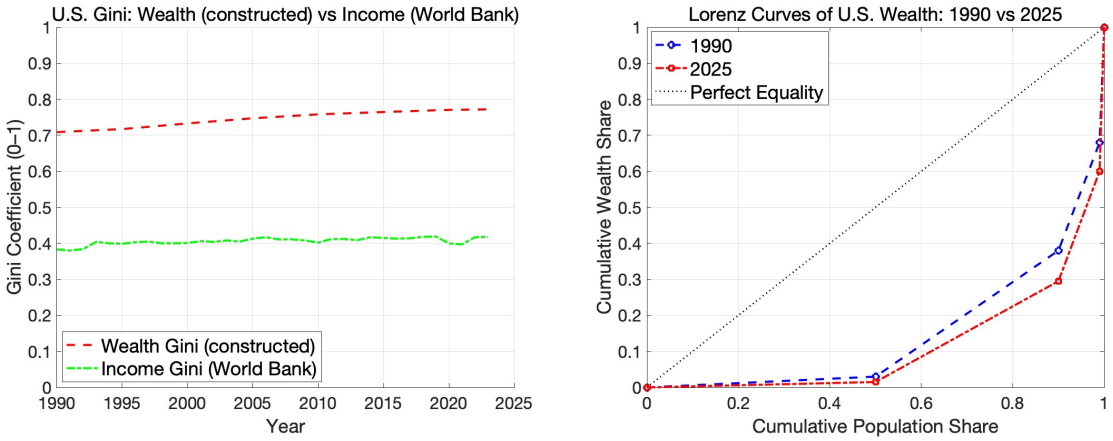
To better capture inequality dynamics, Figure 11 presents the Gini coefficient on wealth as well as on income (left panel) and the Lorenz curves for selected years (right panel), showing that wealth inequality in the U.S. has been increasing over time.

Return Heterogeneity Finally, heterogeneous returns confirm the information-cost channel. Figure 12, 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 the 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, my 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

Figure 11: **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.

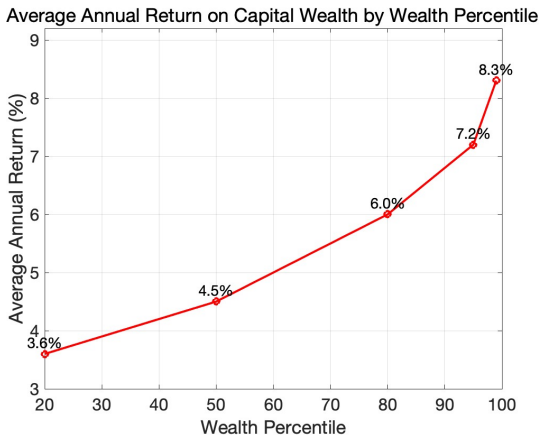


Notes. Income Gini (World Bank) and Wealth Gini (constructed from DFA) by year. Right: Lorenz curves for U.S. wealth in 1990 (blue) and 2025 (red); the 45° line denotes perfect equality. Summary statistics over the sample: Income Gini (World Bank): mean= 38.780, min=34.700, max=41.900. Wealth Gini (constructed): 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 my 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.

canonical Bewley–Aiyagari settings. Precautionary savings models typically cannot produce wealth distributions substantially thicker than the earnings distribution they inherit (DeNardi, Fella and Pardo, 2016; Carroll, Slacalek, Tokuoka and White, 2017; Hubmer, Krusell and Smith, 2016). By contrast, models with return heterogeneity, such as Gabaix, Lasry, Lions and Moll (2016) and Benhabib et al. (2017), are able to reproduce empirically relevant tails. My setting embeds this return heterogeneity channel directly, linking it to financial information frictions and technology.

Figure 12: **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.

F Heterogeneous risk aversion with common initial wealth

This appendix shows that the paper's equilibrium characterization and the incidence results obtain when all investors have the same initial wealth $W_{0j} = W_0$, but have heterogeneous risk aversion ρ_j. The key point is that the mechanism requires heterogeneity in *effective risk tolerance*, not heterogeneity in wealth per se.

F.1 Environment and preferences

All primitives are as in Section 2, except:

- Initial wealth is common: $W_{0j} = W_0$ for all $j \in [0, 1]$.
- Investors differ in risk aversion $\rho_j > 0$, distributed according to G.

Using the standard CARA approximation to CRRA around wealth W_0, investor j's

mean–variance objective at the trading date is

$$\max_{q_{1j}} \ \mathbb{E}_j[W_{2j}] - \frac{\rho_j(W_0)}{2}\mathbb{V}_j(W_{2j}), \qquad \rho_j(W_0) \equiv \frac{\rho_j}{W_0}, \tag{F.1}$$

so that absolute risk tolerance is

$$T_j \equiv \rho_j(W_0)^{-1} = \frac{W_0}{\rho_j}.$$

(Equivalently, one can take T_j as primitive; all results below are stated in terms of T_j.)

The budget constraint remains

$$W_{2j} = r\Big(W_0 - 1_{\text{part},j}F - C(h_j) - p\,q_{1j}\Big) + q_{1j}z, \tag{F.2}$$

with signal $s_j = z + \varepsilon_j$ and cost $C(h) = c_1 h + \frac{c_0}{2}h^2$ as in Section 2.

F.2 Trading demand

Conditional on the trading-date information set $\mathcal{I}_j$, the optimal risky demand is the standard CARA/Normal expression:

$$q_{1j} = T_j\,\frac{\mathbb{E}_j[z] - rp}{\mathbb{V}_j[z]}. \tag{F.3}$$

Relative to the baseline, the only change is that the scale factor is $T_j = W_0/\rho_j$ instead of W_{0j}/ρ.

F.3 Equilibrium aggregation and price

Define the two aggregates (over the endogenous sets of participants and informed investors)

$$t \equiv \int_{j:\,1_{\text{part},j}=1} T_j\,dG, \qquad s \equiv \int_{j:\,h_j>0} T_j h_j\,dG. \tag{F.4}$$

Market clearing with demands (F.3) implies the same linear price form as in the baseline:

$$rp = a + bz - cx,$$

with (a, b, c) functions of (s, t) exactly as in Appendix B.2, and with price precision

$$h_p = \frac{s^2}{\sigma_x^2}, \qquad \bar{h} = h_0 + h_p = h_0 + \frac{s^2}{\sigma_x^2}.$$

Therefore, all objects that depend on the equilibrium price informativeness continue to be pinned down by the single scalar s, while aggregate risk-bearing capacity is summarized by t.

F.4 The marginal value of precision and information choice

Let $D(s, t)$ denote the planning-date marginal gain (in mean–variance units) from one additional unit of private precision. The same algebra as in Appendix B.4 yields:

$$D(s, t) \equiv \frac{\sigma_x^2 + 2st + t^2 h_0}{t^2 \left(h_0 + s^2/\sigma_x^2\right)^2}. \tag{F.5}$$

(As in the baseline, $D(s, t)$ also equals the expected squared Sharpe ratio available to an uninformed participant.)

At $t = 0$, investor j chooses precision $h_j \geq 0$ to maximize the ex-ante gain net of cost:

$$\max_{h_j \geq 0} \ \frac{T_j}{2} D(s, t)\, h_j - c_1 h_j - \frac{c_0}{2} h_j^2.$$

The first-order condition implies the optimal precision schedule

$$h_j = \max\left\{0, \ \frac{T_j D(s, t) - 2c_1}{2c_0}\right\}. \tag{F.6}$$

Thus information acquisition is monotone in risk tolerance T_j (equivalently decreasing in risk aversion ρ_j).

F.5 Participation and information cutoffs (in risk tolerance)

The ex-ante certainty-equivalent gain from entering *without* a private signal is

$$\Delta U_j^{\text{mkt}} = \frac{T_j}{2} D(s, t) - F,$$

so investor j participates iff $\Delta U_j^{\text{mkt}} \geq 0$, i.e.

$$T_j \geq T^{\text{part}} \equiv \frac{2F}{D(s,t)}. \tag{F.7}$$

Similarly, investor j is informed iff $h_j > 0$, which from (F.6) is equivalent to

$$T_j \geq T^{\text{info}} \equiv \frac{2c_1}{D(s,t)}. \tag{F.8}$$

Hence

$$\frac{T^{\text{info}}}{T^{\text{part}}} = \frac{c_1}{F}, \tag{F.9}$$

exactly mirroring the baseline ratio $W_0^{\text{info}}/W_0^{\text{part}} = c_1/F$. When $c_1 > F$, the population partitions into three groups in T-space: nonparticipants ($T < T^{\text{part}}$), uninformed participants ($T^{\text{part}} \leq T < T^{\text{info}}$), and informed participants ($T \geq T^{\text{info}}$).

F.6 Certainty-equivalent wealth schedule and the incidence mapping

With common initial wealth W_0, cross-sectional dispersion in terminal outcomes is entirely endogenous and arises from heterogeneous participation and information choices. The certainty-equivalent terminal wealth of type T can be written as

$$Y(T;\theta) = rW_0 + \left(\frac{T}{2}D(\theta) - F\right)_+ + \left(\frac{(TD(\theta) - 2c_1)^2}{8c_0}\right)_+, \qquad \theta \equiv (F, c_1, c_0), \tag{F.10}$$

where the first "wedge" activates at T^{part} and the second at T^{info}. Equation (F.10) is the exact analogue of equation (13) in the main text, with W_0 replaced by risk tolerance T as the sorting variable. Consequently, the incidence decomposition in Sections 3.1–3.3 goes through verbatim once the primitive cross-sectional distribution is taken to be the distribution of T (equivalently ρ_j) rather than the distribution of W_{0j}.

In particular: (i) $F \downarrow$ shifts the left cutoff T^{part} down directly and (under the same active-information-region conditions used in the baseline) shifts T^{info} up through gen-

eral equilibrium in $D(s,t)$, compressing between-group dispersion; (ii) $c_1 \downarrow$ shifts T^{info} down directly and (again under the same interior conditions) shifts T^{part} up through $D(s,t)$, increasing between-group dispersion; and (iii) $c_0 \downarrow$ steepens the informed wedge in (F.10) and increases within-informed dispersion, while leaving the ratio (F.9) unchanged.

F.7 Why this establishes the robustness claim

The mapping from the baseline to this extension is immediate: in the baseline, $T_j = W_{0j}/\rho$; here, $T_j = W_0/\rho_j$. In both cases, *heterogeneity in* T_j generates heterogeneous values of participation and information, induces sorting, and delivers the same comparative statics and inequality rankings. By contrast, with common $\bar{\rho}$ (CARA with identical risk aversion) and heterogeneous W_{0j}, $T_j \equiv 1/\bar{\rho}$ is constant, so the sorting and distributional mechanisms vanish. This is exactly the sense in which heterogeneity in wealth without heterogeneity in risk tolerance does not deliver the paper's results, while heterogeneity in risk aversion does, even with common wealth.

G Delegation: Details and Proofs

G.1 Mapping Gârleanu and Pedersen (2018) to our setup

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 benchmark); 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).[13] 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 investors' risk tolerances). Aggregate risk tolerance and private precision become

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

$$s = \underbrace{\int_{\text{self-informed}} \rho^{-1}(W_{0j}) h_j\, dV(W_{0j})}_{\text{individual signals}} + \underbrace{T_{\mathcal{M}} h_m}_{\text{manager's pooled signal}}. \tag{G.2}$$

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 benchmark. In particular, the manager's activity enters s additively via $T_{\mathcal{M}} h_m$ (the price coefficients then adjust endogenously just as before).

[13] 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 benchmark corresponds to $\eta = 0$.

G.2 Solving the model with delegation

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

$$\max_{h_m \geq 0} \frac{T_{\mathcal{M}}}{2} D(s,t)\, h_m - c_1 h_m - \frac{c_0}{2} h_m^2$$
$$h_m(T_{\mathcal{M}}; s, t) = \frac{T_{\mathcal{M}} D(s,t) - 2c_1}{2c_0}. \tag{G.3}$$

Therefore a *pool-level* information threshold emerges:

$$T_{\mathcal{M}}^{\text{info}} = \frac{2c_1}{D(s,t)}. \tag{G.4}$$

Pooling creates economies of scale: even if no individual j satisfies the benchmark 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:

$$f \;=\; \frac{C\big(h_m\big)}{M} \;=\; \frac{c_1 h_m + \frac{c_0}{2}(h_m)^2}{M}, \tag{G.5}$$

where $M \equiv \text{mass}(\mathcal{M})$[14] and $D(s,t)$ is the marginal gain per unit precision defined in the benchmark.

Possible extension: Bargaining 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 $f = \frac{C(h_m)}{M} + \eta \frac{JS}{M}$. Our delegation model corresponds to $\eta = 0$.

[14] 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; see Gârleanu and Pedersen (2018) for microfoundations and related delegated-management environments in Cuoco and Kaniel (2011); He and Xiong (2013).

Possible further extension: Full fee menu 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

$$W_{0j} \;\geq\; W_0^{\text{del}} \;\equiv\; \rho \cdot \frac{2\,(f + F_m)}{D(s,t)\, h_m}. \tag{G.6}$$

Compare (G.6) to the benchmark 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 (see Lemma 4) 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

$$\Delta U_j^{self} = \frac{[\rho^{-1}(W_{0j}) D(s,t) - 2c_1]^2}{8c_0}, \tag{G.7}$$

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

$$\Delta U_j^{\text{del}} \equiv \frac{W_{0j}}{\rho}\frac{D}{2}h_m - f - F_m \geq \max\{0,\ \Delta U_j^{self}\}. \tag{G.8}$$

G.3 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 G.8, 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 G.5.

(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 benchmark economy.

(iv) *Consistency (fixed point).* Aggregates are the sums of delegated and self-learning contributions as defined in Equations G.1 and G.2, 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 benchmark; the inner

loop pins down (h_m, M) using (G.3)–(G.5)–(G.6) together with the wealth distribution. Existence and local uniqueness extend from the benchmark economy (Appendix B.6) by continuity. All derivations and details are in Appendix G. 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 this Appendix. We use these objects for classification and comparative statics.

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 benchmark. The fixed point in (s, t) extends the map in Appendix B.6 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 benchmark and is briefly discussed in the main text and Appendix G.2.

G.4 Three lemmas on equilibrium with pooling into funds

Self-learning can be optimal at the top.

Lemma 4 (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. □

Three or four partitions of investor categories: non-participants, uninformed, informed delegators, self-informed

Lemma 5 (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.

Intuition. 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 benchmark economy in Section 2: non-participants, uninformed participants, and self-informed participants.

Proof of Lemma 5 (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 benchmark 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 (G.3) is $h_m > 0$. Throughout, write

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

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

Step 1 (Shape and single crossing). For $W \geq W_{\text{info}}$, $\Delta U^{\text{self}}(W) = \left(\frac{W}{\rho}D - 2c_1\right)^2/(8c_0)$ 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 Lemma 5. □

Delegated information with pooling equilibrium — summary of pool's precision, fees, delegation threshold, price rule and fixed point

Lemma 6 (Delegated information with pooling). *Fix aggregate objects (s,t) and let $D(s,t)$ denote the marginal gain per unit precision defined in the benchmark, and $C(h) = c_1 h_m + \frac{c_0}{2} h_m^2$ the manager's cost. Consider a pool with mass M and aggregate risk tolerance T_M. Statements (i)–(v) hold: i)* Manager's optimal precision and information threshold.

$$h_m = \left[\frac{T_M D(s,t) - 2c_1}{2c_0}\right]_+, \qquad T_M^{info} = \frac{2c_1}{D(s,t)}.$$

(ii) Fees. *Under perfect competition among managers, the per-capita fee equals average cost,*

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

Under a Nash-bargaining variant with manager bargaining weight $\eta \in [0,1]$, the fee is

$$f = \frac{C(h_m)}{M} + \eta\,\frac{JS}{M}, \quad \text{where } JS = \tfrac{c_0}{2} h_m^2.$$

(iii) Delegation decision. *An investor with wealth W_0 and absolute risk tolerance $T = W_0/\rho$ delegates iff*

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

Locally, W_0^{del} increases in F_m, c_1, and c_0.

(iv) Pooling superadditivity (scale). *For interior points ($T \geq T_M^{info}$), aggregate contributed precision is convex in T: for any T_1, T_2,*

$$T_1 h_m(T_1) + T_2 h_m(T_2) \;\leq\; (T_1 + T_2)\, h_m(T_1 + T_2),$$

with strict inequality when $T_1, T_2 > 0$.

(v) Price rule and fixed point. *The delegated block contributes $T_M h_m$ to s, so the price rule remains $r_p = a + bz - cx$ with $b = sc$ and $h_p = s^2/\sigma_x^2$. Existence and local uniqueness of the outer fixed point in (s,t) follow by continuity from the baseline ($D_s < 0$, $D_t < 0$).*

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 benchmark 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. □

G.5 Inequality Effect

Given (s, t, h_m, f), investors sort into: (i) $W_{0j} < W_0^{\text{part}}$ (non-participants); (ii) $W_0^{\text{part}} \leq W_{0j} < \min\{W_0^{\text{del}}, W_0^{\text{info}}\}$ (uninformed); (iii) $W_0^{\text{del}} \leq W_{0j} < W_0^{\text{switch}}$ (delegators); (iv) $W_{0j} \geq W_0^{\text{switch}}$ (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 next.

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}.

Proposition 5 (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 so fees remain unchanged. 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 rises. A cheaper linear information cost boosts both self- learning and delegation, raises s, lowers D, hence decreases participation (pushes* W_0^{part} *right). Some households exit, which* widens *the distribution and increases inequality strongly. At the same time, manager precision* h_m *rises, manager fees fall, and there is more delegation at the left margin. A larger share becomes informed - self or delegated - (*$W_0^{del} = W_0^{info}$ *and* W_0^{switch} *both fall), increasing payoff dispersion within participants. 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, inequality can grow or fall depending on whether mass is in the tails or in the middle.*

4. *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 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 benchmark Theorem 1 (lower information costs increase G).

Under the same mean-preserving normalization as in the benchmark, 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 this 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 Proposition 5. 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).[15] 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 5). 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).

[15] As in the benchmark, 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.

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 (benchmark), 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 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↓ (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:[16]

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

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 proposition.

[16] 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.

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

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

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

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

G.5.1 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 increas-

ing 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 Proposition 5. □

G.5.2 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).[17] 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;[18] terms (i)–(iii) trace the *distribution* channel; term (v) nets out real

[17] Notation as in Appendix D.7.

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

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.[19] *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.

G.6 Quantifying the Inequality Effect

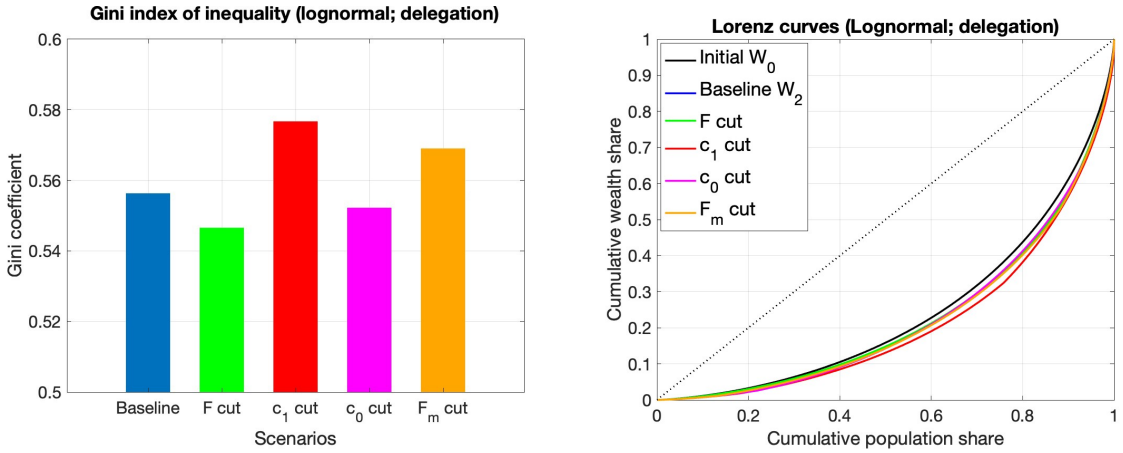
We redo the exercise from the benchmark economy (in the main text) of quantifying the impact of technological progress on inequality in this new setting with delegation and pooling into a fund. The ranking of effects remains virtually unchanged: a cut in participation decreases inequality, while a cut in the marginal information cost c_1 increases inequality unambiguously; a cut in the quadratic information cost decreases inequality under a Lognormal prior, and increases inequality under a Pareto prior.

G.6.1 Under Lognormal

Figure 13 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 16 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 benchmark economy without delegation.

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

Figure 13: **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 8 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 the 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 8: Wealth distribution, thresholds, and inequality across scenarios (lognormal fit), scenario color code stays fixed

Scenario	s	Tm	Ts	D($\times 10^{-6}$)	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	9.3320	1647.23	38.2892	13.783	692.9	4523.6	48633.6	82.80%	23.80%	0%	0.568

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 benchmark economy without delegation: reducing F delivers the strongest compression of inequality, reducing the Gini coefficient; reducing c_1 raises the Gini coefficient 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 coefficient, with the c_0 effect theoretically ambiguous and quantitatively modest in my calibration.[20]

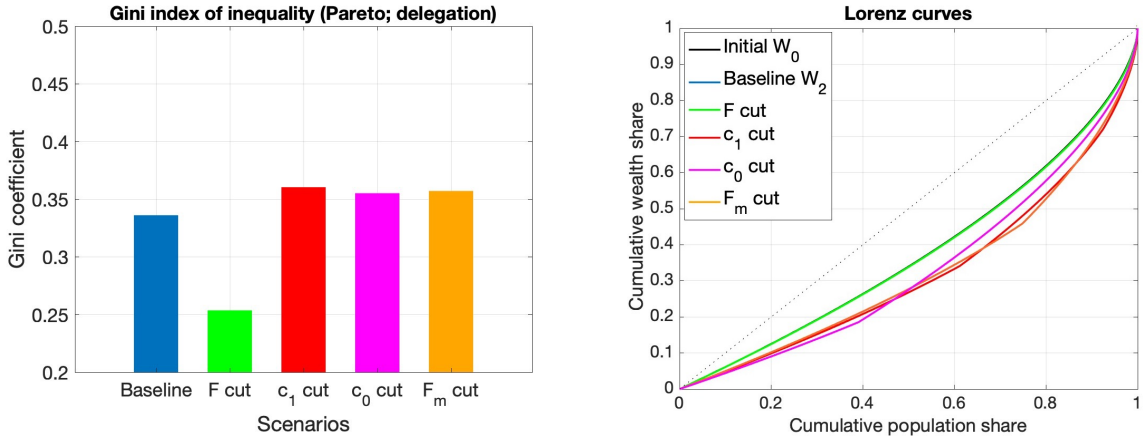
G.6.2 Under Pareto

Figure 14 plots the Gini index and the Lorenz curves in respond to each cost cut in the model with manager delegation and pooling. Similarly, Table 9 reports the four cutoffs and inequality across the five scenarios.

Table 9 shows that the ordering of scenarios is unchanged relative to the benchmark model without delegation: cutting the participation cost F delivers the largest *decline* in inequality (Gini coefficient 0.336 $\rightarrow$ 0.254), while lowering information costs raises inequality, and an F_m cut is neutral in my 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

[20] In the table, the c_0 cut yields a small decline in the Gini coefficient 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 14: **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.

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 coefficient 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 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 coefficient 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 coefficient even with delegation.

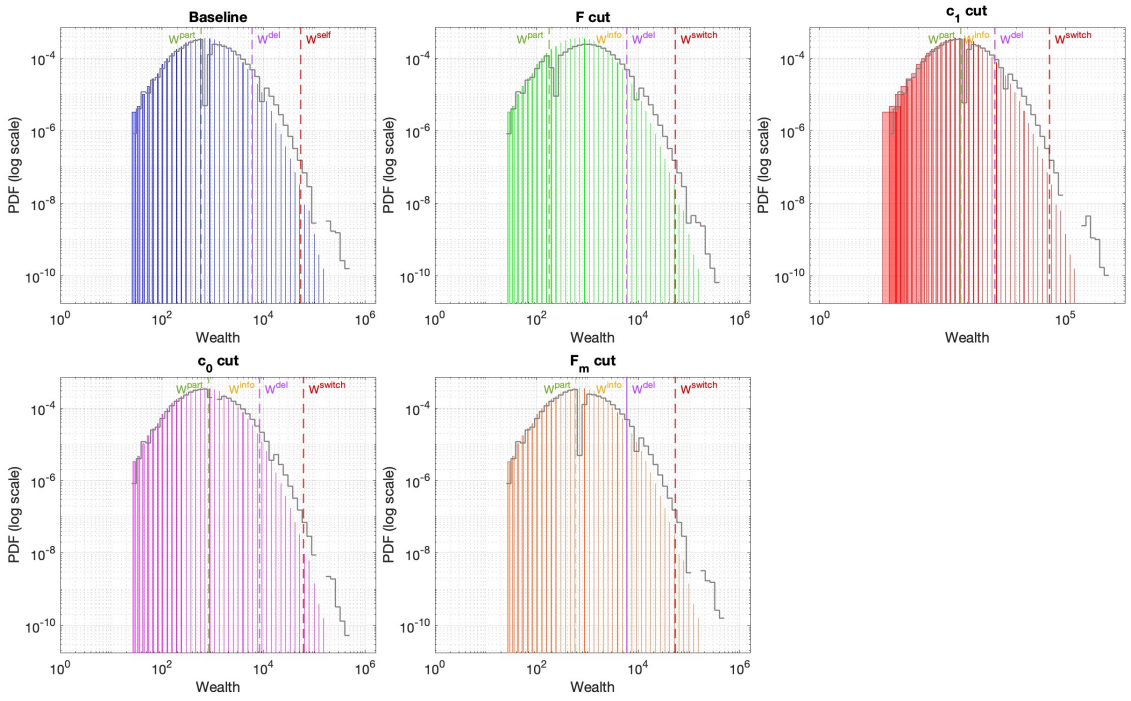
Table 9: Wealth distribution, thresholds, and inequality across scenarios (Pareto fit), scenario color code stays fixed

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	3.15	29.3777	0.50	1433.22	161.03	371.28	5397.32	41	13.4	0%	0.356

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.

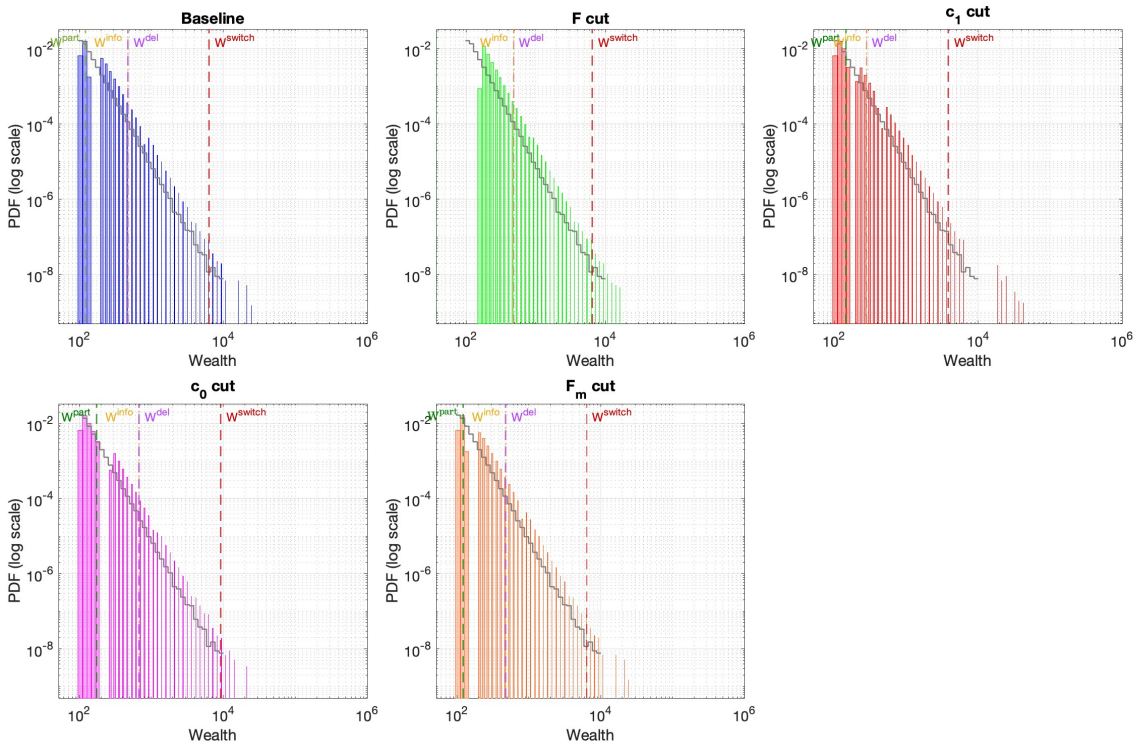
G.7 Additional tables and figures for delegating extension

Figure 15: **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 16: **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.

H Discussion: Policy-relevant implications of the incidence decomposition

This paper is a positive theory of *distributional incidence*: it characterizes how technology parameters $\theta = (F, c_1, c_0)$ map into the cross-section of certainty-equivalent (CE) terminal wealth and hence into Lorenz/Gini-type inequality objects. The model is *not* designed to deliver a complete welfare analysis: it abstracts from several objects that a full normative evaluation would require (e.g., the welfare of noise traders and possible externalities of private information production). Accordingly, the discussion below is intentionally brief and focuses on what the incidence decomposition implies for how one should *think about* (and empirically evaluate) different classes of financial-technology interventions.[21]

H.1 Technology shocks are distinct distributional levers

Theorem 1 implies that "financial technology" is not a single lever. Different innovations move different wedges in the CE schedule, and therefore differ in the *incidence* of gains:

1. **Entry / participation technology** ($F \downarrow$) primarily operates on the extensive margin of market access. By shifting mass across the participation cutoff W_0^{part}, it compresses the gap between nonparticipants and participants, and is therefore the model's cleanest channel for Lorenz-improving changes (lower measured inequality) in CE wealth.
2. **Information-adoption technology** ($c_1 \downarrow$) primarily operates on the extensive margin of becoming informed. By shifting mass across the uninformed–informed boundary W_0^{info}, it reallocates CE gains toward the informed segment and tends to thicken the upper tail among participants. In the model this is the cleanest channel

[21] This is also why the main text emphasizes incidence rankings and internal consistency checks rather than policy prescriptions. :contentReference[oaicite:0]index=0

for Lorenz-worsening changes (higher measured inequality).

3. **Information-intensity technology** ($c_0 \downarrow$) primarily operates on the intensive margin of precision *within* the informed group. This lever is therefore the natural mechanism for changes in *within-informed dispersion*, even when the aggregate Lorenz/Gini response is modest or ambiguous due to general-equilibrium feedback through price informativeness.

A practical implication is that policy discussions framed as "lowering frictions" are incomplete: *which* friction matters because it determines whether the first-order incidence runs through between-group reallocation (across participation/information cutoffs) or through within-informed dispersion (precision heterogeneity).

H.2 Inclusion does not hinge on the participation margin

A second implication is that the distributional forces associated with information technology do not require a large participation margin. Even holding participation fixed (or in high-participation baselines), the adoption and intensity levers continue to reshape the within-participants distribution by reallocating CE gains across informed versus uninformed traders and by altering within-informed dispersion. Put differently, "democratization" in the form of broader market access need not prevent inequality from rising if the technological change primarily lowers c_1 or c_0 and thus concentrates gains among informed investors or increases within-informed dispersion.

H.3 What to measure when evaluating technology interventions

Because the model's incidence decomposition has sharp mapping from technology levers to observable intermediate objects, an empirical or policy evaluation should, at a minimum, report diagnostics that correspond to the decomposition:

1. **Extensive-margin diagnostics:** participation rates, and proxies for the informed share (or for informed activity), together with the implied movements in the participation and adoption cutoffs.
2. **Between-group wedges:** differences in risk-adjusted returns (or CE proxies) between participants and nonparticipants, and between informed and uninformed participants.
3. **Within-informed dispersion:** cross-sectional dispersion in proxies for information intensity/precision among informed traders, and conditional inequality measures computed within the informed group.

Reporting these objects alongside standard price-efficiency outcomes (e.g., informativeness) makes it possible to distinguish "efficiency improvements with progressive incidence" (more consistent with $F \downarrow$) from "efficiency improvements with regressive incidence" (more consistent with $c_1 \downarrow$ and, for within-group dispersion, $c_0 \downarrow$).

H.4 Delegation and packaged access

Delegation can be interpreted through the same lens: is a *technology for reallocating information access*, potentially transforming a convex, wealth-sensitive information wedge into a more proportional exposure to information-driven gains. In incidence terms, the extension speaks to whether a given expansion in "being informed" manifests primarily as (i) a between-group reallocation (uninformed $\rightarrow$ informed) or (ii) an increase in within-informed dispersion. This is why the extension is best read as a mechanism that can *reshape the decomposition*, rather than as a separate normative policy analysis.