The Organization of Information Production in Financial Markets

Jun Aoyagi*

HKUST

August 1, 2026

Abstract

Recent advances in AI are reshaping the organization of financial information production. This paper develops a theory of how technological progress reshapes information production through the interaction between labor markets and financial markets. Existing views emphasize that AI decentralizes information production by expanding individual capabilities. I show that this captures only one channel: technological progress can intensify competition in financial markets and eliminate the profitability of organized information production, generating a fundamentally different form of decentralization. Technological progress can concentrate or decentralize information production depending on the stage of technological development, with implications for liquidity, welfare, and technology design.

JEL classification: D83, G14, G23, L23, O33

Keywords: information production, technology, organization, externality

*junaoyagi@ust.hk. 5005 LSK Business School, Clear Water Bay, Hong Kong University of Science and Technology.

1 Introduction

Recent advances in artificial intelligence (AI) and other information technologies have dramatically increased individuals' ability to acquire and process financial information. A growing literature argues that these technologies democratize access to sophisticated information processing and reduce barriers to active trading (e.g., Chang, Dong, Margin, and Zhou, 2026), leading to the view that talented individuals and small investment firms can increasingly perform tasks that previously required large research organizations. Similarly, the recent industry reports (e.g., Bloomberg, 2026; Forbs, 2026) describe how AI-powered investment firms are replacing large analyst teams with a handful of researchers and AI systems, reconfiguring the organization of information production. Existing theories explain how technological progress affects information acquisition and market efficiency, but provide little guidance on how it reshapes the *organization* of information production. This naturally raises the question of whether technological progress will fundamentally decentralize financial information production. More generally, why does technological progress sometimes reshape the organization of information production rather than simply making it more productive?

This paper develops a model of endogenous information production through financial agents' occupational choices and demonstrates that the relationship between technological progress and the organization of information production is fundamentally non-monotonic.[1] The model consists of a continuum of heterogeneous agents who differ in their information-processing skills. An agent's information production capacity is jointly determined by her skill and the technology level common across all agents. Each agent first chooses between establishing a trading firm as an owner-manager or serving as a financial analyst. Firms and analysts participate in the labor market, where trading firms organize information production by hiring analysts and combining their information into a private signal about a risky asset. Firms then participate in the financial market by incurring a fixed participation cost, and the asset price aggregates information through the noisy order flow.[2] In equilibrium,

[1]The main argument of the paper focuses on general-purpose information technologies, such as the Internet and generative AI, that are broadly accessible and enhance agents' ability to generate an information edge in financial markets. I abstract from special-purpose technologies, such as high-frequency trading technologies, that benefit only specific agents.

[2]The financial market in my model combines the Kyle (1985) price-formation mechanism with competition among informed traders in the spirit of Holden and Subrahmanyam (1992), while incorporating

occupational choice is characterized by a skill cutoff: high-skilled agents become managers and organize information production, whereas low-skilled agents become analysts and provide their information-processing capacity to trading firms. The endogenous interaction between the analyst labor market and the financial market determines both the organization of information production and market outcomes.

Technological progress has two opposing effects on the organization of information production. First, by improving each agent's information-processing capability, it raises the return to establishing a trading firm on her own. This induces intermediate-skilled agents to establish their own trading firms instead of working as analysts, thereby reducing the supply of analysts available to incumbent firms. I refer to this force as the *capability effect*. Second, technological progress intensifies competition among informed trading firms. Because firms fail to internalize the information revelation through their own trades (Grossman and Stiglitz, 1980), aggregate information production erodes informational rents in the financial market, weakening firms' incentives to organize information production. I refer to this force as the *competition effect*. The interaction between these two forces generates a non-monotonic evolution of the organization of information production.

At relatively low levels of technological development, information revelation is not substantial, and the capability effect dominates the competition effect. Even intermediate-skilled agents, who would normally lack the ability to organize information production to trade on their own, can benefit from technological innovation and enter the market as trading firms (managers). Consequently, the supply of analysts declines, and each firm is forced to operate with a relatively small information-production team. Organized information production therefore becomes more decentralized: a large number of trading firms, each supported by a small analyst team, prevail in the economy. This pattern is referred to as *capability-driven decentralization* and corresponds to the conventional view that technology broadens participation of a wide range of agents in the financial market by expanding individual capabilities of information production.

As technology continues to improve, however, competition among informed trading firms intensifies and progressively erodes informational rents, thereby inducing intermediate-skilled

quadratic trading costs, following Vives (2008), to prevent complete information revelation.

managers to exit and return to the supply side of the labor market. Organized information production therefore experiences *competition-driven concentration*: only the most skilled trading firms remain active, each employing a large analyst team. Rather than broadening participation, further technological progress now reinforces concentration through the competition effect. This phase is consistent with patterns occasionally observed in financial markets, where large firms increasingly dominate talent acquisition, making it harder for smaller firms to compete (Business Insider, 2025).

Beyond a critical level of technological development, however, the competition effect becomes sufficiently strong that the hiring equilibrium ceases to exist. As informational rents are substantially eroded by competition, establishing a trading firm is no longer profitable once the fixed participation cost is taken into account.[3] Consequently, the managers' participation constraint fails, implying that the competitive wage required to sustain the analyst labor market would be negative. Organized information production therefore collapses together with the analyst labor market. The economy instead transitions to a *decentralized equilibrium*, where information production and trading are carried out by stand-alone managers without organized analyst teams. Once the labor market disappears, low-skilled informed agents no longer produce information under organized firms, mitigating information revelation and the competition externality. The resulting recovery in informational rents restores the profitability of stand-alone trading for sufficiently skilled agents, even though organized information production remains unsustainable. Information production therefore shifts discontinuously from analyst teams to independent traders, giving rise to a second, fundamentally different form of decentralization. This transition is referred to as *the competition-driven decentralization*. Unlike capability-driven decentralization, this force arises not because individuals become more capable, but because intensified competition eliminates the profitability of organized information production.

The model delivers three main implications. First, it identifies two fundamentally different channels through which technological progress decentralizes information production. The conventional capability effect reflects expanding individual information-processing capa-

[3]Even at zero wages, an incumbent manager would still benefit from hiring additional analysts. Rather, the remaining informational rents are too small to justify becoming a manager in the first place.

bility, whereas the competition effect arises because intensified competition erodes informational rents and endogenously collapses organized information production. Although both mechanisms generate smaller organizations and more independent traders, they have opposite implications for informational rents and market liquidity, providing a natural empirical strategy to distinguish them. Second, the model predicts that the organization of information production evolves non-monotonically with technological progress: decentralization at early stages of technological progress is followed by concentration before organized information production eventually collapses. This result provides a unified theoretical framework for understanding why technological advances can lead either to decentralized information production or to increased concentration among a small number of dominant financial firms, as highlighted in recent concerns about market concentration (e.g., Aquilina, Budish, and O'Neill, 2021). Finally, the model shows that organized information production need not improve the welfare of agents. The organized information production encourages low-skilled agents to produce information under trading firms, while they would have stayed inactive in the absence of the labor market for analysts. The additional information contributes to intensifying competition in financial markets and erodes the very informational rents that information producers seek to exploit. This finding highlights a novel trade-off between organizational information productivity and the profitability of information, suggesting an important tension in the design of information technologies and policy implications for technology adoption in financial markets.

A growing literature studies how advances in AI and other information technologies affect knowledge production and information processing (Brynjolfsson, Li, and Raymond, 2023; Korinek, 2023; Chang, Dong, Margin, and Zhou, 2026). This paper contributes to this literature by studying how such technologies endogenously determine the organization of information production. In finance, Farboodi and Veldkamp (2020) show that information technology reshapes the type of information traders choose to produce. More broadly, several studies analyze information aggregation through organization (e.g., Becker and Murphy, 1992; Zandt, 1999; Garicano, 2000; Garicano and Rossi-Hansberg, 2006). Closest to my setting, Stein (2002) shows that financial firms choose between hierarchical and decentralized information production depending on whether information can be credibly communicated up a chain of command. In these papers, the limiting force of organization is an internal cost, such as communication costs, so that technological changes affect the size of organizations in a monotonic manner. In contrast, my model embeds a competitive financial market that disciplines the expansion of organization, demonstrating non-monotonic and discontinuous evolution of organized information production.

A separate literature studies how competition among traders erodes the value of information in asset markets. Grossman and Stiglitz (1980) establish the seminal result on the interplay between information acquisition by individual traders and competition. Stein (2009) extends this logic to "crowded trades," showing the self-defeating nature of competing traders, while Glode, Green, and Lowery (2012) show that competitive over-investment in costly expertise is fully dissipated in equilibrium and can trigger a breakdown in trade when adverse selection is severe. More broadly, a large literature studies information acquisition under competitive rational-expectations equilibria (e.g., Verrecchia, 1982; Admati, 1985; Veldkamp, 2006). Unlike this literature, in which agents individually choose how much information to acquire, my model features information production through a rival input (human capital), showing that technological progress reshapes not only how much information is produced but also how information production is organized.

2 Model

Consider a three-period economy with a continuum of *agents*. In the first period ($t = 0$), agents with heterogeneous skills exert costly effort to acquire financial knowledge. In the second period ($t = 1$), agents sort themselves into one of two occupations: establishing a *trading firm* or working as a *financial analyst*. Each trading firm can be seen as an investment fund operated by a single owner-manager, who hires analysts to organize information production and makes trading decisions.[4] In the third period ($t = 2$), trading firms invest in the financial market based on information generated by analysts on staff. A single risky asset is traded in the financial market whose payoff, v, follows a normal distribution with

[4] Throughout the paper, I use "trading firm" and "(owner-)manager" interchangeably whenever no confusion arises.

mean $\bar{v}$ and variance σ_v^2 and is realized at the end of $t = 2$.

2.1 Skill and Knowledge

Each agent has information-production capacity, η, which determines the precision with which an agent can assess the magnitude of mispricing in the asset (specified below). The value of η depends on idiosyncratic skills and common information technology, as specified by the following information-production function:

$$\eta_z \equiv \alpha z \theta, \tag{2.1}$$

where $z \in [0, 1]$ and $\theta \in \{0, 1\}$ represent *information-processing skill* and *financial knowledge*, respectively, and they are heterogeneous across agents. Parameter $\alpha \in [\alpha_0, 1]$ is the information productivity common across all agents, capturing the level of information technology that helps agents collect and process financial information, such as generative AI, personal computers, and the Internet.[5]

At the beginning of $t = 0$, agents draw information-processing skill z from a uniform distribution, $z \sim U[0, 1]$. This skill can be interpreted as the agent's inherent computational ability, which is determined exogenously and varies across agents. Upon observing z, each agent chooses whether to exert effort to acquire financial knowledge by incurring a fixed non-pecuniary effort cost $\epsilon > 0$. The acquisition of financial knowledge $\theta \in \{0, 1\}$ depends on an agent's effort choice: by exerting effort, she learns financial knowledge and obtains $\theta = 1$ with probability one, but if she does not exert effort, she obtains $\theta = 1$ only with probability $\phi < 1$. The result of knowledge acquisition is realized in the end of $t = 0$. This effort choice can be seen as studying for CFA or pursuing a finance degree: passing the exam or completing the program allows an agent to apply her information-processing skill z to analyze the financial market.

In what follows, z denotes an agent's *type*, and an agent who acquires financial knowledge ($\theta = 1$) is referred to as an *informed* agent, whereas an agent who does not ($\theta = 0$) is referred

[5]The lower bound of the technology level, $\alpha_0 > 0$, is imposed to ensure the existence of an equilibrium with financial market trading.

to as an *uninformed* agent. Information about (z, θ) stays private to each agent. Also, the effort cost is positive but assumed to be diminishing, $\epsilon \to 0$. This assumption greatly simplifies the model without affecting the qualitative results.

2.2 Labor Market

In the second period ($t = 1$), each agent chooses her role in the financial market—either establishing a trading firm or serving as an analyst—and participates in the labor market. For simplicity, the following analyses assume that only informed agents can establish a trading firm, while Appendix A demonstrates that this is the optimal choice endogenously arises in the equilibrium.

If a type-z agent establishes a firm, she may hire measure q_z of analysts. This hiring requires the payment of wage w per analyst, and the firm also incurs a quadratic managerial cost, $m(q_z) = \frac{\gamma}{2} q_z^2$ with $\gamma > 0$. This managerial cost is introduced to avoid unlimited demand for analysts. In reality, such a cost arises because hiring and deploying a larger workforce entails increasing organizational costs, shaping m increasing and convex in q_z. As specified below, each trading firm generates a private signal about the asset payoff by utilizing the aggregate information-production capacity of informed analysts together with the owner-manager's own capacity. I refer to such a hierarchical arrangement as the *organization* of information production, and the size of the organization is measured by q_z. Note that a trading firm can always choose not to hire analysts ($q_z = 0$) and instead engage in independent information production by using only the manager's own information-processing capacity.

Alternatively, an agent may choose to be an analyst. Such an agent receives wage income w and is hired by a trading firm through a competitive labor market.[6] If the analyst is informed, then she contributes to the information production by the firm as $\eta_z > 0$; if she is uninformed, $\eta_z = 0$, and she has no contribution.[7] In what follows, I denote the sets of agents who establish trading firms as owner-managers and financial analysts as M and A,

[6] An analyst may trade the asset through independent information production, but such an agent is classified as a trading firm with a stand-alone manager according to the definition.

[7] Uninformed analysts can have positive contribution to information production by assuming that their knowledge level is $\theta = \theta_0$ with $0 < \theta_0 < 1$. This alternative assumption does not change the main results of the paper, while complicating the analyses.

respectively.

The labor market for financial analysts is competitive and anonymous. The competitiveness means that the wage w is determined by the market clearing condition for the labor market, and the anonymity implies that firms choose their demand for analysts and hire them without knowing which analysts are informed. Importantly, analysts on the supply side of the labor market receives the pooling wage w regardless of their financial knowledge, so that w can be seen as the upfront signing bonus that is paid prior to performing the actual information production tasks.

2.3 Information

Suppose that a trading firm of type z has hired measure q_z of analysts. It utilizes the aggregate information-production capacity, including its own, to generate a private signal about the asset payoff. In particular, I define the firm's aggregate information-production capacity by

$$\tau_z \equiv \eta_z + q_z \mu_A, \tag{2.2}$$

where

$$\mu_A \equiv \mathrm{E}[\eta_z | z \in A] \tag{2.3}$$

represents the average information-production capacity of informed analysts on staff. The first term of τ_z captures the firm manager's own contribution to the total capacity, and the second term represents the total contribution of hired analysts. By leveraging τ_z, the firm observes a private signal about the asset payoff, $s_z = v + \xi_z$, where the informativeness of this signal is directly proportional to τ_z. In particular, the noise term ξ_z follows a normal distribution with mean zero and variance $\sigma_z^2 \equiv \sigma_v^2 \frac{1-\tau_z}{\tau_z}$, implying that the signal informativeness becomes

$$\frac{\mathrm{Var}[v]}{\mathrm{Var}[s_z]} = \tau_z. \tag{2.4}$$

A larger aggregate information-production capacity τ_z reduces the noise variance of the signal s_z and yields a more precise private information to the firm. The firm observes the signal s_z at the beginning of the third period ($t = 3$) and makes its investment decision based on s_z.

2.4 Financial Market

The financial market in the last period ($t = 2$) has three types of participants: trading firms, a noise trader, and a competitive market maker.[8] I assume that each trading firm incurs a fixed participation cost $\kappa > 0$ to trade in the financial market, capturing fixed expenditures associated with establishing research infrastructure, following compliance, and maintaining operational setup (e.g., Vissing-Jørgensen, 2002).

Upon participating in the financial market, a trading firm of type z places a market order for $x_z \in \mathbb{R}$ units of the asset, where the firm faces a quadratic transaction cost $c(x_z) = \frac{\rho}{2}x_z^2$ with $\rho > 0$. This cost can be seen as the risk aversion or inventory pressure and ensures that the optimal investment is finite. The noise trader places an exogenous and random market order for u units of the asset, where u follows a normal distribution with mean zero and variance σ_u^2. The market maker observes the total order flow, $y \equiv \int_{z \in M} x_z dz + u$, and sets the execution price of the asset, p. Due to competition, the price is determined by the semi-strong efficient level:

$$p = \mathrm{E}[v|y]. \tag{2.5}$$

After all orders are executed, the asset payoff v is realized, and each firm earns trading profit $\pi_z \equiv (v - p)x_z - \frac{\rho}{2}x_z^2$.

2.5 Utility and Equilibrium

An agent of type z obtains the following utility by establishing a trading firm, hiring q_z analysts, and trading x_z units of the asset:

$$U_{z,M} = \pi_z - wq_z - \frac{\gamma}{2}q_z^2 - \kappa,$$

where the second and the third terms are the wage payment and the managerial costs, and the last term captures the fixed participation cost of the financial market. If an agent

[8]Following the convention in the market microstructure literature, I assume that there are competitively many market makers off the equilibrium path, and only one of them executes incoming orders on the equilibrium path. This setting ensures the competitive asset price because, if the active market maker deviates from the competitive price, then other market makers would quote a better price.

becomes an analyst, in contrast, her utility stems from the wage payment,

$$U_{z,A} = w.$$

The utility of the noise trader directly reflects the price impact of her order u and is formalized later in Section 5. The market maker breaks even by trading at the competitive price.

The equilibrium of the model is defined as follows:

Definition 1. The equilibrium of the model consists of $(A, M, \{x_z, q_z\}_{z\in M},\ p,\ w)$ and individual agents' effort choice, such that; (i) each agent chooses her effort to maximize her expected utility, either $\mathrm{E}[U_{z,A}]$ or $\mathrm{E}[U_{z,M}]$, incorporating the effort cost ϵ; (ii) the sets of analysts and trading firm, A and M, are determined by each agent's occupational choice that maximizes her expected utility, $\mathrm{E}[U_{z,A}]$ or $\mathrm{E}[U_{z,M}]$; (iii) a trading firm of type z chooses the trading quantity x_z to maximize the expected trading profit π_z conditional on signal s_z; (iv) it chooses the demand for analysts q_z to maximize the expected utility $\mathrm{E}[U_{z,M}]$; (v) the asset price p is determined by the efficiency condition in (2.5); and (vi) the wage w is determined by the market-clearing condition of the labor market.

Finally, I impose the following parameter conditions:

Assumption 2. $\sigma_v > \max\left\{\sqrt{\frac{\rho\kappa}{\alpha_0}}, \frac{4\rho\sigma_u}{\phi}\right\}$.

The assumption guarantees the existence of at least one type of equilibrium (specified below) with active trading firms by ensuring that the value of acquiring and trading on private information is sufficiently high.

3 Equilibrium

3.1 Financial Market

The model is solved by backward induction. I guess and later verify that a trading firm of type z in the financial market ($t = 2$) takes the following linear trading strategy.

$$x_z = \beta\tau_z(s_z - \bar{v}), \tag{3.1}$$

where β is an endogenous constant determined in the equilibrium and is referred to as the trading intensity. In equation (3.1), $\tau_z(s_z - \bar{v})$ represents the *ex-ante* mispricing in the asset: the market maker's prior expectation is $\mathrm{E}[v] = \bar{v}$, and she would quote $\bar{v}$ without observing the order flow, while the informed trading firm's conditional expectation is $\mathrm{E}[v|s_z] = \bar{v} + \tau_z(s_z - \bar{v})$. Thus, the difference between them, $\tau_z(s_z - \bar{v})$, can be seen as the asset's mispricing, and its magnitude is scaled by the firm's aggregate information-production capacity, τ_z. The strategy in (3.1) implies that the firm exploits this mispricing with intensity β.

According to the trading strategy in (3.1), the aggregate order flow from trading firms is defined as

$$X \equiv \int_{z \in M} x_z dz = \beta\tau(v - \bar{v}), \tag{3.2}$$

where

$$\tau \equiv \int_{z \in M} \tau_z dz. \tag{3.3}$$

As τ captures the *economy-wide* information-production capacity that all active agents devote to analyze the risky asset, it is referred to as the *aggregate information quality* in the following discussion. From the market maker's perspective, the aggregate order flow, $y = X + u$, is linear in v.[9] The standard filtering leads to the following efficient price.[10]

$$p = \bar{v} + \lambda y, \tag{3.4}$$

where

$$\lambda = \frac{\beta\tau\sigma_v^2}{\beta^2\tau^2\sigma_v^2 + \sigma_u^2}. \tag{3.5}$$

The coefficient λ represents the price impact of order flow and, following Kyle (1985), λ is used as the measure of illiquidity: a high λ implies an illiquid financial market.

Since the hiring and the participation costs are sunk at the trading stage, a trading firm

[9]The condition for the exact law of large numbers is satisfied in this model, as $\int_{z \in M} \sigma_z^2 dz < \infty$, and idiosyncratic noise terms $\{\xi_z\}_{z \in M}$ in equation (3.2) average out in the cross section, implying $\int_{z \in M} \xi_z dz = 0$ almost surely.

[10]It holds that $p = \mathrm{E}[v|y] = \bar{v} + \frac{\mathrm{Cov}[v,y]}{\mathrm{Var}[y]}(y - \mathrm{E}[y])$, where $\lambda = \frac{\mathrm{Cov}[v,y]}{\mathrm{Var}[y]}$ yields the result in (3.5).

of type z chooses its trading strategy x_z to maximize the expected net trading profit:

$$\max_{x_z} \mathrm{E}[(v-p)x_z|s_z] - \frac{\rho}{2}x_z^2, \tag{3.6}$$

where the firm incorporates the pricing rule in (3.4) and the aggregate order flow (3.2), while its price impact is zero since each firm's measure is infinitesimal. The first-order condition of (3.6) yields the following optimal strategy.

$$x_z = \frac{1-\lambda\beta\tau}{\rho}\tau_z(s_z - \bar{v}),$$

suggesting that the guess in (3.1) is correct with the trading intensity

$$\beta = \frac{1-\lambda\beta\tau}{\rho}. \tag{3.7}$$

Solving equations (3.5) and (3.7) yields β and λ as functions of τ:

Lemma 3. *(i) The optimal trading strategy of a trading firm, β, is determined by the unique solution to the following fixed-point problem and is monotonically decreasing in the aggregate information quality τ.*

$$\rho\beta = \frac{\sigma_u^2}{\beta^2\tau^2\sigma_v^2 + \sigma_u^2}. \tag{3.8}$$

(ii) By applying the solution β, the price impact, λ, is determined by (3.5) and takes a single-peaked reaction to τ.

The price impact exhibits a hump-shaped curve against τ. Intuitively, τ captures the total information advantage of active trading firms, and an increase in τ imposes a larger adverse selection cost on the market maker, thereby inducing her to set a larger price impact. However, as τ becomes sufficiently large, the order flow reveals excessive information about v from the perspective of individual trading firms, because each firm ignores the information-revelation externality generated by its own trades. Anticipating this information revelation, the market maker places less weight on the order flow when inferring the asset value, leading to a lower price impact. While λ takes a hump-shaped reaction to τ, the price impounds information v with coefficient $\beta\lambda\tau$, which is monotonically increasing in τ (see [3.5]), suggesting that the price sensitivity of v increases with τ. To avoid moving the price too much, each manager trades less aggressively by lowering β when τ becomes larger.[11]

Furthermore, applying (3.5) and (3.7), the expected trading profit of the type-z trading firm, conditional on s_z, is computed as

$$\pi_z = \rho\beta^2\tau_z^2(s_z - \bar{v})^2.$$

Intuitively, the firm's informational advantage is $\tau_z(s_z - \bar{v})$, and each trading unit expects to earn this profit margin. Since the firm scales the trading strategy x_z by its informational advantage as well, the total expected profit becomes quadratic in $\tau_z(s_z - \bar{v})$. Also, by taking the unconditional expectation of π_z,

$$\bar{\pi}_z = \mathrm{E}[\pi_z] = \rho\sigma_v^2\beta^2\tau_z. \tag{3.9}$$

Therefore, the firm's *ex-ante* expected trading profit (prior to observing s_z) becomes proportional to the aggregate information-production capacity τ_z of the firm's organization.

3.2 Labor Market

This subsection explores how the organization of information production is shaped through agents' occupational choices in the labor market.

[11] In the Kyle (1985) model, the insider trades as a monopolist in the financial market by fully incorporating the information revelation via her trading activity. In this model, in contrast, informed managers do not internalize this effect, because individual managers are infinitesimal and have no price impact. As a result, they reveal too much information, causing the negative reaction of λ to τ.

3.2.1 Trading Firm's Demand for Analysts

Given the expected trading profit $\bar{\pi}_z$ in (3.9), a type-z trading firm chooses its demand for analysts by solving the following problem.[12]

$$\max_{q_z \geq 0} \bar{\pi}_z - wq_z - \frac{\gamma}{2} q_z^2.$$

Since the firm is infinitesimal, it does not internalize the impact of its analyst demand q_z on the aggregate information quality τ. Noting that the firm's information-production capacity is $\tau_z = \alpha z + q_z \mu_A$, the first-order condition yields the following optimal analyst demand:

$$q^* = \max\left\{\frac{\rho\sigma_v^2\beta^2\mu_A - w}{\gamma}, 0\right\}. \tag{3.10}$$

The marginal profit of hiring analysts, μ_A, is independent of the firm's type z, making the optimal demand q^* common across all firms. Equation (3.10) suggests the possibility of two types of equilibria: one with active hiring, $q^* > 0$, and the other with no labor market activity, $q^* = 0$. In what follows, I focus on the active hiring equilibrium and leave the analysis of the non-hiring equilibrium for Subsection 3.3.

In the hiring equilibrium, the aggregate demand for analysts is given by

$$Q = \int_{z \in M} \left(\frac{\rho\sigma_v^2\beta^2\mu_A - w}{\gamma}\right) dz. \tag{3.11}$$

The type-z firm expects to obtain the following indirect utility by hiring q^* analysts and paying the fixed participation cost:

$$\mathrm{E}[U_{z,M}] = \alpha\rho\sigma_v^2\beta^2 z + \frac{(\rho\sigma_v^2\beta^2\mu_A - w)^2}{2\gamma} - \kappa. \tag{3.12}$$

The first term represents the expected net trading profit arising from the firm's own information-processing skill z, and the second term is that arising from hiring analysts, net of the wage

[12] Note that the market-participation constraint is redundant. Any agent who chooses to become a manager must earn a higher payoff net of the participation cost than the analyst wage w. Since $w \geq 0$, it follows that the manager's payoff before paying the participation cost must exceed the cost itself, implying that market participation is always optimal.

payment and the managerial cost.

3.2.2 Occupational Choice and Effort

A type-z informed agent compares the expected utility of a trading firm $\mathrm{E}[U_{z,M}]$ in (3.12) with the analysts' wage income w: she establishes a trading firm if and only if $\mathrm{E}[U_{z,M}] \geq w$.[13] Since $\mathrm{E}[U_{z,M}]$ in equation (3.12) linearly increases with an agent skill z, this comparison is characterized by a cutoff, z^*. In particular, informed agents with relatively high skills, $z \geq z^*$, choose to establish a firm, while uninformed agents with $z \geq z^*$, as well as agents with relatively low skills, $z < z^*$, become a financial analyst.

Moreover, the cutoff z^* characterizes the effort choice as well. Namely, agents with $z < z^*$ would not found a trading firm, even if they acquire financial knowledge, and end up receiving the pooling wage w regardless of their effort choice. Since the effort choice does not affect their expected utility, while exerting effort is costly, they optimally choose not to exert effort. In contrast, the occupational choice of agents with relatively high skill, $z \geq z^*$, depends on the result of their knowledge acquisition: if an agent becomes informed, she chooses to establish a trading firm, while she becomes an analyst otherwise. As $\mathrm{E}[U_{z,M}] \geq w$ holds for these agents, the marginal benefit of effort is positive. As the effort cost is diminishing ($\epsilon \to 0$), they optimally choose to exert effort.

Lemma 4. *There is a cutoff z^* as a unique solution to the following indifference condition:*

$$\alpha\rho\sigma_v^2\beta^2 z^* + \frac{\left(\rho\sigma_v^2\beta^2\mu_A - w\right)^2}{2\gamma} = w + \kappa. \tag{3.13}$$

Agents with type $z \geq z^*$ exert effort and establish a trading firm, while those who with type $z < z^*$ do not exert effort and become an analyst.

Lemma 4 implies that the set of analysts is expressed as $A = \{z|z < z^*\}$, while the set of trading firms is $M = \{z|z \geq z^*\}$. Hence, the organization of information production in the model is characterized by the cutoff z^*: agents with relatively low skills ($z \in A$) are hired by trading firms operated by managers with relatively high skills ($z \in M$). Analysts

[13] An agent is assumed to establish a firm when she is indifferent between the two occupational role, i.e., $\bar{U}_{z,M} = w$. This tie-breaking rule is, however, innocuous to the results.

produce information about the risky asset under the supervision of these trading firms. As the cutoff z^* increases, a larger mass of agents become analysts, and information production is concentrated to a smaller set of high-skilled trading firms. Conversely, a decline in z^* induces more agents to establish a firm by exerting effort. The information production is more dispersed across a large number of firms characterized by a wide range of skills.

3.2.3 Labor Market Clearing

According to Lemma 4, the aggregate analyst supply is given by $S \equiv \int_0^{z^*} dz = z^*$, and equation (3.11) leads to the following labor market clearing condition:

$$\frac{\rho\sigma_v^2\beta^2\mu_A - w}{\gamma}(1-z^*) = z^*. \tag{3.14}$$

In the equilibrium, the wage and the cutoff adjust so that individual firms' demand for analysts become $q^* = \frac{z^*}{1-z^*}$. Namely, mass $1-z^*$ of trading firms demand q^* of analysts each, while the supply of analysts amounts to mass z^*.

3.2.4 Analyst Quality and Information Production

On the supply side, informed and uninformed analysts are pooled together with mass ϕ and $1-\phi$, and the anonymity of the labor market leads to the quality uncertainty form the demand-side perspective. However, the average information productivity of hiring analysts, μ_A in equation (2.3), becomes deterministic due to the law of large numbers:

$$\mu_A = \mathrm{E}\left[\alpha z\theta | z < z^*\right] = \frac{\alpha\phi z^*}{2}. \tag{3.15}$$

As the cutoff z^* increases, agents with relatively high skills choose to be an analyst, thereby improving the average quality of the analyst pool in the labor market.

Furthermore, according to (3.3), the cutoff rule in Lemma 4 provides the explicit formula for the aggregate information quality in the financial market:

$$\tau = \int_{z\in M}(\alpha z + q^*\mu_A)dz = \alpha\frac{1-(1-\phi)z^{*2}}{2}, \tag{3.16}$$

where $\int_{z\in M} q^* dz = z^*$ holds due to the labor market clearing condition. Interestingly, the aggregate information quality τ is decreasing in z^*, showing the opposite behavior compared to the analysts' information productivity μ_A. When z^* increases, a larger fraction of agents remains in the "no effort" regime, where expected productivity by these agents is only ϕz rather than z. The information production loss from these additional agents dominates the compositional effect in the labor market that raises the conditional average μ_A, causing a decline in τ. In what follows, I denote the total information quality as $\tau(z^*)$ and the trading intensity $\beta(z^*) = \beta(\tau(z^*))$ to specify their dependence on z^*, where Lemma 3 and equation (3.16) imply that $\beta(z^*)$ is monotonically increasing in z^*.

By incorporating (3.14)–(3.16), the hiring equilibrium is characterized by the following:

Proposition 5 (Hiring Equilibrium)**.** *There is at most one hiring equilibrium characterized by the cutoff z^*. The cutoff is the unique solution to the following condition*

$$\alpha\rho\sigma_v^2\beta^2(z^*)z^*\left(1-\frac{\phi}{2}\right)+\gamma\frac{z^*\left(1-\frac{z^*}{2}\right)}{(1-z^*)^2}=\kappa, \tag{3.17}$$

where $\beta(z^*)$ is the unique solution to (3.8).

The LHS of equation (3.17) represents the expected net trading profit of a trading firm of type z^*, after incorporating the hiring cost and the opportunity cost of serving as an analyst (i.e., the wage income), while κ on the RHS is the participation cost of the financial market. As the LHS is monotonically increasing in z^*, condition (3.17) ensures a unique solution z^*.[14]

3.3 Decentralized Equilibrium

The solution z^* to (3.17) constitutes the hiring equilibrium only if the wage level and the analyst demand induced by z^* are positive. As elaborated in Section 4, such an equilibrium may not exist. Even in the absence of the hiring equilibrium, however, the other equilibrium with no labor market activity always exists :

[14]The LHS of (3.17) becomes zero at $z^* = 0$ and explodes to positive infinity as $z^* \to 1$. Hence, condition (3.17) always has a unique solution in $z^* \in (0, 1)$.

Proposition 6 (Decentralized Equilibrium)**.** *There always exists an equilibrium with no labor market activity. It is characterized by the cutoff z_d^*, such that, agents with $z \geq z_d^*$ exert effort and establish trading firms, while those who with $z < z_d^*$ do not exert effort and stay inactive in the subsequent periods. A firm's trading strategy, $\beta_d(z_d^*)$, is the unique solution to*

$$\rho\beta = \frac{\sigma_u^2}{\beta^2 \tau_d^2(z_d^*)\sigma_v^2 + \sigma_u^2}$$

with

$$\tau_d(z_d^*) = \frac{\alpha(1 - z_d^{*2})}{2}. \tag{3.18}$$

The cutoff z_d^* is determined by the following indifference condition:

$$\alpha\rho\sigma_v^2\beta_d^2(z_d^*)z_d^* = \kappa. \tag{3.19}$$

As in the hiring equilibrium, agents with relatively high skills ($z \geq z_d^*$) exert effort and become informed, while those with low skills ($z < z_d^*$) do not. Low-skilled agents stay inactive in the subsequent periods regardless of the result of knowledge acquisition because no wage is paid in the labor market ($w = 0$). The absence of the labor market also implies that high-skilled informed agents do not hire an analyst; instead they join the financial market as stand-alone traders. The equilibrium in Proposition 6 is referred to as the *decentralized* equilibrium, because no hierarchical structure for information production emerges, and all information is produced and utilized in trading through stand-alone individual managers operating trading firms.

The aggregate information quality, τ_d, and thus the equilibrium trading strategy of firms are slightly different from those in the hiring equilibrium (τ and $\beta(z^*)$ in equation [3.16] and Lemma 3), due to the absence of hiring. This difference captures the loss of information in decentralized equilibrium, $\tau_d(z) < \tau(z)$ at the same level of z, and leads to more intensive trading of individual firms, $\beta_d(z) > \beta(z)$. Due to the absence of the labor market, agents who do not exert effort but obtain financial knowledge must stay inactive in the decentralized equilibrium, whereas they would have produced information as analysts if the hiring equilibrium existed and the labor market was active.

As the decentralized equilibrium always exists, the existence of a hiring equilibrium necessarily implies equilibrium multiplicity, a feature common to models of adverse selection in asset markets (e.g., Kurlat, 2013; Fukui, 2018). Because the multiplicity of equilibria is not the focus of this study, I assume that the economy coordinates on the hiring equilibrium whenever it coexists with the decentralized one. This selection is supported by Wilson (1980), Stiglitz and Weiss (1981), and Kurlat (2013), who argue that a no-trade (no-hiring) equilibrium is not robust to buyers' (employers') incentive to unilaterally raise prices (wages) to attract a higher-quality pool, and accordingly focus on the highest price-quantity equilibrium whenever it exists. In my model, the same argument holds: hiring a small mass of analysts is beneficial for a trading firm when $w = 0$, while low-skilled agents who do not exert effort but obtain financial knowledge are willing to serve as analysts whenever $w > 0$. Thus, such deviation is supported as a hiring equilibrium whenever z^* as the solution to (3.13) guarantees $w(z^*) > 0$. When z^* induces $w(z^*) \leq 0$, in contrast, the decentralized equilibrium is a unique and stable outcome.

4 Organization of Information Production

This section explores the impact of technology improvement (increases in α) on the organization of information production. I first focus on the behavior of the hiring equilibrium in Subsection 4.1. In Subsection 4.2, its existence condition is specified, and the phase transition between the hiring and the decentralized equilibrium is characterized.

4.1 Hiring Equilibrium

4.1.1 Information Externality

The indifference condition in equation (3.13) suggests that an increase in the information-production technology α influences the cutoff z^* by changing the expected *profit margin* of information production, which is defined as

$$R(z^*, \alpha) \equiv \alpha\rho\sigma_v^2\beta^2(\tau(z^*, \alpha)), \tag{4.1}$$

where I denote $\beta(z^*) = \beta(\tau(z^*, \alpha))$ to underline the dependence of the trading strategy on the technology level α. Namely, a trading firm of type z leverages its own skill to earn $R(z^*, \alpha)$. On top of that, each unit of information-production capacity acquired through hiring analysts also earns $\rho\sigma_v^2\beta^2\mu_A = \frac{\phi}{2}R(z^*, \alpha)$. The profit margin is common across all firms, as the price impact in the financial market is determined by the total information quality τ. The following lemma shows that α affects R through multiple channels, leading to the non-monotonic reaction.

Lemma 7. *With the cutoff z^* being fixed, a trading firm's expected profit margin takes a single-peaked curve with respect to α. There is a unique tipping point α^*, such that, $\frac{\partial R(\alpha, z^*)}{\partial \alpha} \gtreqless 0 \Leftrightarrow \alpha^* \gtreqless \alpha$, where α^* is defined as a unique solution to (B.4) in Appendix B.*

Firstly, α has *the capability effect*: it directly increases the profit margin R by expanding the information-production capability of a firm's organization, as captured by the first coefficient in (4.1). The expansion of information production, however, has an indirect negative effect on R, which is referred to as *the competition effect*. Namely, an increase in α leads to a higher aggregate information quality τ, and the order flow $y = X + u$ reveals substantial information about v to the market maker. Consequently, each firms must hold back from trading aggressively by lowering β and accept a smaller profit margin. If a single trading firm trades as a monopolistic informed trader (e.g., Kyle, 1985) and perfectly internalizes the information revelation through the trading activity, the firm could adjust its trading strategy β so that these two competing effects completely offset. In my model, by contrast, competitively many firms trade in the financial market and cannot perfectly internalize information revelation of their trading activity. Due to the information externality, the negative indirect effect of α can dominate its positive direct effect, leading to the ambiguous reaction of R to changes in α.

4.1.2 Information Production

Reflecting the non-monotonic impact of the technology level α on the profit margin of each firm, the cutoff z^* exhibits the following reaction to changes in α.

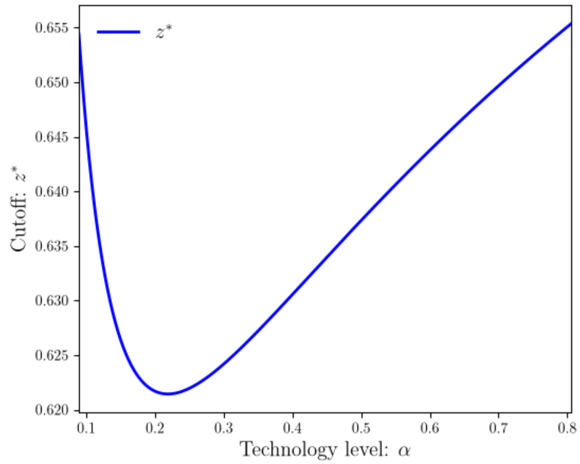

Figure 1: The hiring-equilibrium cutoff z^*

Note: This figure plots the cutoff in the hiring equilibrium using the following parameter values: $\phi = 0.2, \kappa = 2.5, \gamma = 0.12, \rho = 0.1, \sigma_v^2 = 7.0$, and $\sigma_u^2 = 1.0$.

Proposition 8. *The cutoff z^* exhibits a U-shaped reaction to increases in α. With the tipping point α^* defined in Lemma 7, it holds that $\frac{dz^*}{d\alpha} \gtreqless 0 \Leftrightarrow \alpha \gtreqless \alpha^*$.*

Figure 1 visualizes the result in Proposition 8. If the informational profit margin R increases with α, all firms experience an exogenous upward shift in their expected utility, and agents with skill levels just below the cutoff z^* are now willing to establish a trading firm, thereby lowering the cutoff in condition (3.13). If, in contrast, R declines with α, then agents with skill levels slightly hither than the cutoff cease to operate trading firms and stop exerting effort, so that the cutoff z^* increases. As a consequence, the cutoff z^* exhibits a U-shaped curve against α, as opposed to the profit margin R.

This result implies a non-monotonic evolution of the organization of information production. When the technology level is relatively low ($\alpha < \alpha^*$), a technological progress causes decentralization of information production: more agents, including those with relatively low skills, become managers and seek to establish trading firms. As each firm can hire a smaller mass of analysts ($q^* = \frac{z^*}{1-z^*}$ holds in the hiring equilibrium), the economy consists of a large number of small-sized information production organizations. This evolution is triggered by the expansion of individual information-production capability and is referred to as *capability-driven decentralization.*

In contrast, when technology is matured ($\alpha > \alpha^*$), a further increase in the technology

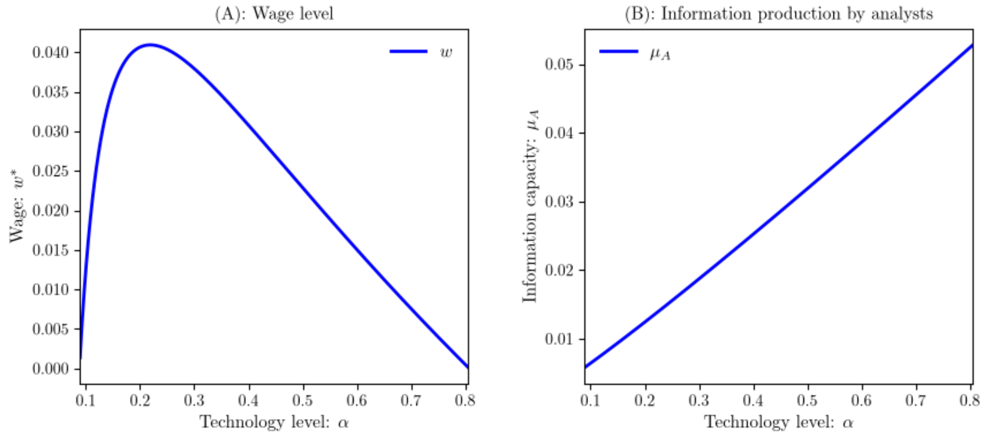

Figure 2: The labor market

Note: This figure plots the wage level, w, and the average information production by analysts, $\mu_A = \frac{\alpha\phi z^*}{2}$, in the hiring equilibrium. The parameter setting is the same as that used in Figure 1.

level induces *competition-driven concentration* of organized information production. Namely, an increase in α lowers the profit margin R and raises the cutoff z^*, pushing a marginal trading firm to stop operating and to switch to the analyst job. This phase exhibits a reversal toward a concentrated structure, where a small number of high-skilled trading firms organize a large information-production organization.

4.1.3 Labor Market

The labor market for financial analysts exhibits the following reaction to the technological progress:

Proposition 9. *(i) The analysts wage, w, takes a single-peaked curve with respect to the technology level α, that is, $\frac{dw}{d\alpha} \gtreqless 0 \Leftrightarrow \alpha^* \gtreqless \alpha$, where the tipping point α^* is the same as that in Proposition 8.*
(ii) The average information capacity of analysts, μ_A in equation (3.15), is monotonically increasing in α.

Figure 2 illustrates the results in Proposition 9. The non-monotonic reaction of the wage is a direct consequence of the U-shaped cutoff z^*. When an increase in the technology level α lowers the cutoff z^* (i.e., when $\alpha < \alpha^*$), a larger mass of trading firms participate in the demand side of the labor market, while the analyst supply shrinks. Since the labor

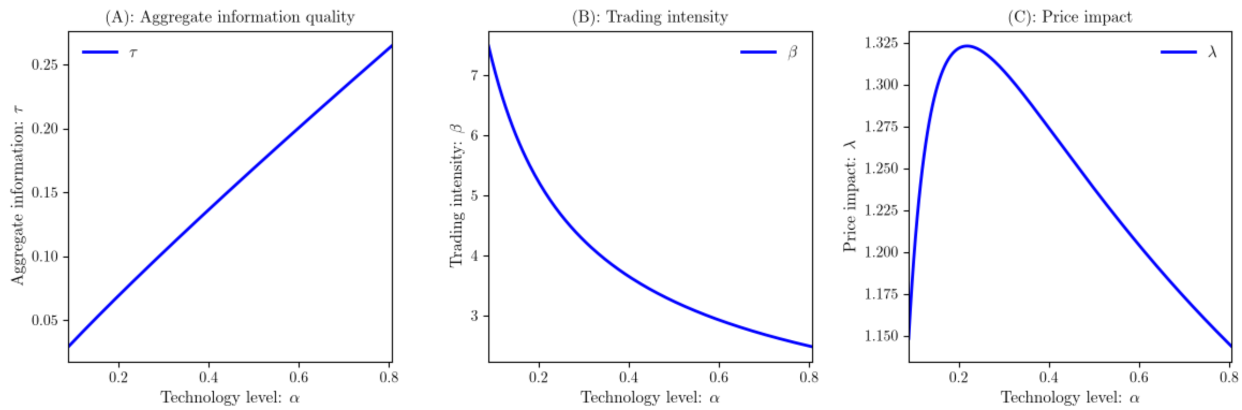

Figure 3: The financial market

Note: This figure plots the aggregate information quality, τ, the trading intensity, β, and the price impact, λ, in the hiring equilibrium. The parameter setting is the same as that used in Figure 1.

market must clear, the intensive margin of the analyst demand must shrink, as captured by $q^* = \frac{z^*}{1-z^*}$. To support this reaction, the corresponding wage level increases. As a relatively high technology level ($\alpha > \alpha^*$) induces the opposite argument, the wage level takes a single-peaked curve against α.

Along with this change in the composition of the labor market, the average information productivity of analysts on the supply side monotonically increases, as Panel B in Figure 2 illustrates. Firstly, a higher technology level α directly expands the information-production capacity of each analyst, thereby increasing μ_A. Secondly, the ambiguous reaction of z^* may or may not improve the average skill of analysts in the labor market. However, this second channel is an indirect effect of α through the labor market and cannot dominate its first direct effect, resulting in the monotonically increasing information productivity of analysts.

4.1.4 Financial Market

The financial market exhibits the following reaction to the technological progress:

Proposition 10. *(i) The aggregate information quality τ, defined by equation (3.16), is monotonically increasing in the technology level α.*
(ii) Firms' trading intensity β is monotonically decreasing in α.
(iii) The price impact λ takes a single-peaked curve against α, that is, $\frac{d\lambda}{d\alpha} \gtreqless 0 \Leftrightarrow \alpha^{**} \gtreqless \alpha$, where the tipping point α^{**} is defined in Appendix B.

Although the cutoff z^* and the organization of information production exhibit non-monotonic reactions to technological progress, the total information production, measured by τ, monotonically increases with α. The mechanism behind this reaction is the same as that supporting the monotonic reaction of the average information production of analysts μ_A in Proposition 9.

Responding to this heightened information production τ, the trading intensity β monotonically declines, as each trading firm seeks to avoid revealing too much information about the asset fundamentals. Due to the information externality mentioned in Subsection 4.1.1, the price impact exhibits a single-peaked curve against α.

4.2 Collapse of the Organization

As mentioned in Subsection 3.3 the hiring equilibrium may disappear when the solution z^* of the indifference condition (3.13) cannot support a positive wage, $w > 0$. The following analyses explore what technology levels support the hiring equilibrium and, when it collapses, how the economy transitions from the hiring to the decentralized equilibrium. To focus on interesting scenarios where the hiring equilibrium exists under a certain range of α, the following analyses assume that condition (C.2) in Appendix C holds.[15]

4.2.1 Technology Thresholds

Due to the hump-shaped reaction of the wage w to the technology level α (Proposition 9), w can be negative both when α is sufficiently small and large, leading to the following existence condition of the hiring equilibrium.

Proposition 11. *(i) There are two thresholds of the technology level, α_l and α_h ($> \alpha_l$), defined by the following equation: for $j \in \{l, h\}$,*

$$\alpha_j \beta^2(\tau(z_0, \alpha_j)) = \frac{2\gamma}{\phi\rho\sigma_v^2(1 - z_0)}, \tag{4.2}$$

with $z_0 \equiv \frac{\phi\kappa+\gamma-\sqrt{\gamma\kappa\phi^2/2+\gamma^2}}{\phi\kappa+2\gamma(1-\phi/4)}$. The hiring equilibrium exists if, and only if, $\alpha \in (\alpha_l, \alpha_h)$. If

[15] If this condition is violated, the decentralized equilibrium becomes unique for all parameter values, and the model is unable to analyze the organization of information production.

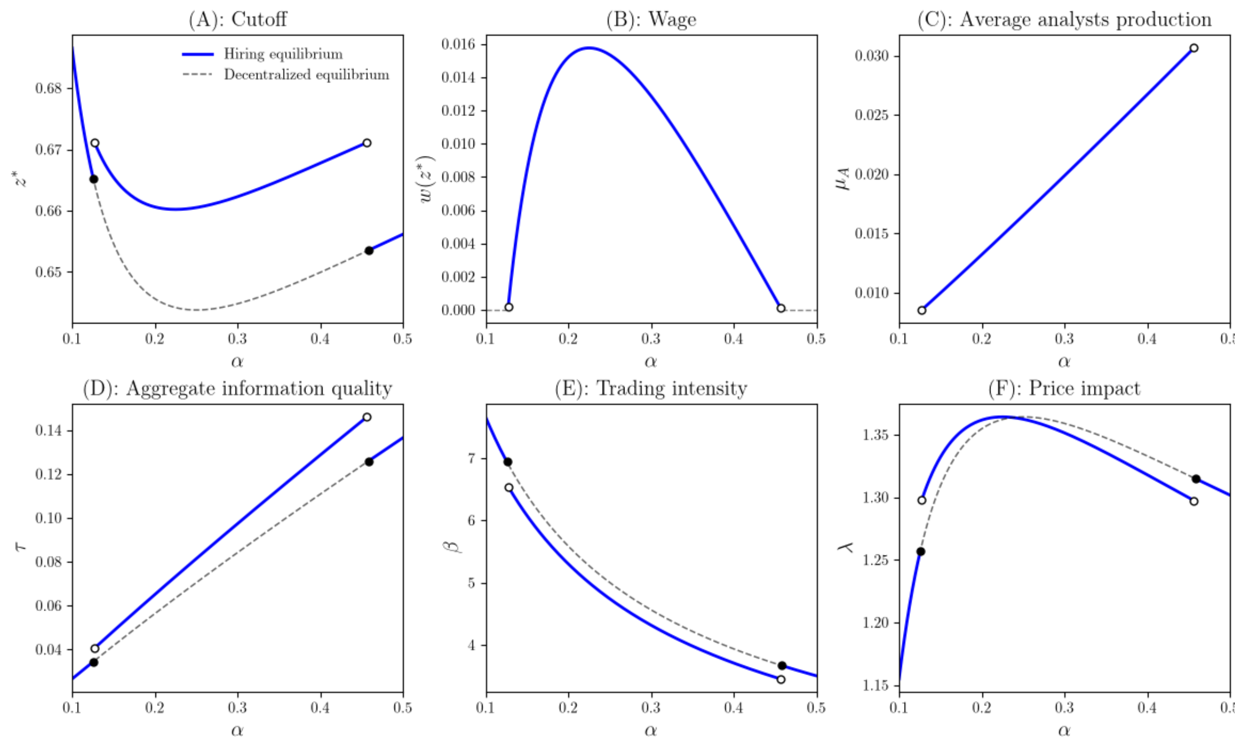

Figure 4: Phase transition

Note: This figure plots the equilibrium variables for the entire range of $\alpha \in [\alpha_0, 1]$. The parameter setting is the same as that used in Figure 1.

$\alpha \notin (\alpha_l, \alpha_h)$, *then the decentralized equilibrium is a unique outcome.*

(ii) At the phase transition from the hiring equilibrium to the decentralized equilibrium, the cutoff experiences a discontinuous shift from z^* *to* z_d^**, where it is a downward jump* ($z^* > z_d^*$) *if, and only if,*

$$\frac{\beta(\tau_d(z_0, \alpha_j))}{\beta(\tau(z_0, \alpha_j))} > \frac{\kappa\phi}{2\gamma} \frac{\gamma\left(1 - \frac{\phi}{2}\right) + \sqrt{\gamma\kappa\phi^2/2 + \gamma^2}}{\phi\kappa + 2\gamma\left(1 - \frac{\phi}{4}\right)}. \tag{4.3}$$

Otherwise, the transition from the hiring to the decentralized equilibrium involves an upward jump in the cutoff.

Figure 4 plots the equilibrium variables against the entire range of α. The figure focuses on the parameter values that satisfy condition (4.3), generating $z^* > z_d^*$, while Figure 6 in Appendix C illustrates the phase transition with $z_d^* > z^*$.

Proposition 11 attests that the hiring equilibrium exists only when the technology level α is intermediate. Firstly, the collapse of the hiring equilibrium under low levels of α is a direct consequence of insufficient information production. In such a situation, agents' information-production capacity tends to be small (given z), and the profit margin in the financial market becomes too small to sustain the organized information production. Secondly, even high levels of α can eliminate the hiring equilibrium. This is because excessive information production intensifies competition among firms and the information externality (Lemma 7), eroding profit margins despite the superior technology. Consequently, the model exhibits *competition-driven decentralization* of information production. That is, firms are discouraged to hire analysts for information production, and the wage level is suppressed (Proposition 9), so that the economy transitions to the decentralized equilibrium, where all information production and trading are carried out by stand-alone managers without forming analyst teams. This result implies that the hierarchical structure of information production ceases to exists not only when the technology level is low but also when it is sufficiently high. Moreover, as discussed below, the phase transition involves abrupt changes in the labor market and the organization of information production, which in turn feeds back to the financial market variables.

4.2.2 Discontinuous Phase Transition

When the economy moves from the hiring to the decentralized equilibrium, both at $\alpha = \alpha_l$ and $\alpha = \alpha_h$, the cutoff z^* exhibits a discontinuous jump. For example, a downward jump in z^* at the threshold α implies that a strictly positive mass of agents with skill levels around z^* start exerting effort for knowledge acquisition.

This result arises from three competing effects of the labor market breakdown. Firstly, the labor market in the hiring equilibrium serves as the outside option for low-skilled informed agents ($z < z^*$): they do not exert effort, but in case that they successfully acquire financial knowledge with probability ϕ, they can participate in the labor market as analysts and receive wage $w > 0$. This backup option, however, does not exist in the decentralized equilibrium, increasing the relative value of being a stand-alone manager and encouraging agents to exert effort. However, the impact of this effect is not significant at the thresholds, as the wage level and the value of the outside option are arbitrary close to zero, i.e., $w(z^*) \to 0$ at $\alpha = \alpha_l$ and $\alpha = \alpha_h$.

Secondly, the absence of the outside option for informed low-skilled agents also implies the loss of information: in the hiring equilibrium, they contribute to information production under a trading firm, while this information is lost in the decentralized equilibrium. This effect is captured by $\tau(z^*, \alpha) > \tau_d(z_d^*, \alpha)$ in equations (3.16) and (3.18). Hence, in the decentralized equilibrium, trading firms bring less information into the financial market, competition among them is mitigated, and the price impact becomes lower, allowing every surviving firms trades more aggressively for a given skill level, as captured by $\beta(\tau_d(z^*, \alpha)) > \beta(\tau(z_d^*, \alpha))$. Since individual firms' profit improves as the economy switches from the hiring equilibrium to the decentralized equilibrium, the cutoff z^* declines, and more agents start exerting effort. This effect is captured by the LHS of condition (4.3). As the strict inequality $\tau(z^*, \alpha) > \tau_d(z_d^*, \alpha)$ holds even when $w(z^*) \to 0$, this effect remains significant at the thresholds, contributing to the discontinuity in the phase transitions.

Thirdly, firms in the decentralized equilibrium cannot hire analysts to expand information production even when hiring analysts is profitable. Thus, a transition from the hiring to the decentralized equilibrium reduces firms' expected profits and raises the bar for knowledge acquisition z^*. This channel also contributes to the discontinuity at the phase transition, as the marginal benefit of hiring analysts remains positive ($q^* > 0$) even when $w = 0$, while they cannot hire due to the absence of the analyst supply in the decentralized equilibrium. Overall, condition (4.3) reflects the competition between the second negative effect and the third positive effect of the labor market breakdown on the cutoff z^*.

4.2.3 Financial Market

As shown in the lower panels of Figure 4, the financial market variables experience an abrupt change at the phase transition. As the decline in the cutoff z^* at the transition from the hiring to the decentralized equilibrium causes a discontinuous decline in the total information quality (τ), it leads trading firms to trade more aggressively in a discontinuous manner.

Interestingly, the reaction of the price impact (and thus market illiquidity) is asymmetric between the lower ($\alpha = \alpha_l$) and the upper ($\alpha = \alpha_h$) thresholds. At the lower threshold, the aggregate information quality is relatively low, and the price impact responds positively when the order flow becomes marginally more informative. In contrast, at the upper threshold, the order flow reveals too much information about the asset fundamentals. The market maker seeks to avoid incorporating abundant information and reduces the price impact when the information quality marginally improves. This asymmetric reaction of the market maker across high- and low-information quality regimes causes the price impact to jump in opposite directions at the two thresholds, even though the cutoff z^* jumps in the same direction.

4.3 Implications

4.3.1 Technological Progress and the Organization of Information Production

Technological innovations, such as recent advances in generative AI, have substantially lowered the cost of processing complex financial information. A growing literature argues that these technologies improve the ability of retail investors and small investment firms to analyze public information and narrow the informational gap with institutional investors. For example, recent evidence in Chang, Dong, Margin, and Zhou (2026) suggests that AI improves retail investors' ability to process earnings-call information and increases their trading profitability, particularly for information-intensive firms. Industry discussions similarly emphasize that AI is "leveling the playing field" by enabling individuals to perform analytical tasks that previously required large research teams (Forbes, 2026).

Our model shows that this intuition captures only one of two fundamentally different mechanisms through which technological progress can decentralize information production. The conventional view is supported by capability-driven decentralization in Proposition 8. Namely, the cutoff z^* in the hiring equilibrium declines in response to an increase in α when $\alpha \leq \alpha^*$ due to the capability effect. At relatively low levels of technological development, improvements in information technology indeed encourage decentralization by expanding the set of individuals who can profitably establish their own investment teams. As information-processing technology improves, agents with increasingly lower skill levels choose to trade as managers over being employed as analysts, leading to an industry populated by a larger number of smaller traders and organizations. This phase is consistent with the existing views mentioned above: technological progress increases the supply of entrepreneurial information producers through the individual capability effect.

Our model identifies a second, fundamentally different source of decentralization. When technological progress becomes sufficiently matured, the aggregate production of information intensifies competition in financial markets and compresses informational rents. As the profitability of active trading declines, firms no longer find it profitable to sustain large analyst teams. Consequently, the analyst labor market collapses, and many former analysts become independent traders (Proposition 11). Importantly, through this competition-driven decentralization, they switch their role not because they are substantially more productive than before, but because organized information production is no longer economically viable.

Although both mechanisms generate a more decentralized information-production structure, they have opposite implications for the profitability of information. In capability-driven decentralization, changes in the organization structure is accompanied by rising informational rents and expanding the mass of trading firms. In competition-driven decentralization, the changes in information production reflect declining informational rents, shrinking analyst employment, and the collapse of organized information production.

This distinction also yields a testable empirical implication. While both mechanisms predict that technological progress improves financial information quality (τ) and induces informed traders to trade less aggressively (β), they imply opposite predictions for price impact (λ). Under capability-driven decentralization, price impact increases with technological progress, whereas under competition-driven decentralization, it decreases. Hence, organizational structure alone cannot distinguish the two mechanisms, since both predict smaller organizations and more independent traders. Instead, our model predicts that the joint evolution of organizational structure and price impact provides a simple empirical criterion for identifying which mechanism drives decentralization.

4.3.2 Concentration versus Decentralization

Recent advances in general-purpose technologies have also bring about concerns that financial information production may become increasingly concentrated (Sherwin, 1981; Autor, Dorn, Katz, Patterson, and Van Reenen, 2020).[16] As information technologies improve, highly

[16]The literature on superstar firms and technology-driven market concentration predicts that technological progress disproportionately benefits the most productive organizations.

skilled investment firms are expected to exploit these technologies more effectively than their competitors, allowing them to expand their informational advantage and attract a disproportionate share of trading profits (Aquilina, Budish, and O'Neill, 2021).

Competition-driven concentration in our model partly confirms this intuition: Proposition 8 demonstrates that an increase in α induces a higher cutoff z^* when $\alpha \geq \alpha^*$. As technology improves beyond the initial expansion phase ($\alpha \leq \alpha^*$), organized information production becomes increasingly concentrated among a smaller number of highly skilled investment firms. Intensifying competition erodes the profitability of marginal firms, causing less productive organizations to exit and leaving information production concentrated in a few large firms with extensive analyst teams. However, this trend does not continue indefinitely. Once technology becomes sufficiently advanced, further improvements reduce the profitability of organized information production too low to sustain the labor market activity. Beyond this phase transition, information production becomes decentralized where all active agents produce information as stand-alone managers.

This generates a novel empirical prediction. Existing theories typically imply a monotonic relationship between technological progress and market concentration. By contrast, our model predicts a non-monotonic pattern: technological progress increases concentration in information production when the technology level is relatively high, but beyond a threshold, a further technological progress triggers a rapid organizational reversal toward decentralized production. Consequently, a decline in concentration need not indicate weaker information technology or lower information quality. Instead, it may reflect a market in which information has become so excessive that organized information production is no longer economically sustainable.

5 Welfare and Policy Implications

Under the hiring equilibrium, an agent's ex-ante expected surplus is defined by

$$U_h = \int_{z^*}^{1} \mathrm{E}[U_{z,M}]dz + \int_0^{z^*} w dz,$$

where the subscript "h" represents the hiring equilibrium. Namely, if she draws a high skill ($z \geq z^*$), she exerts effort and becomes a firm, earning $\mathrm{E}[U_{z,M}]$ in expectation. If, in contrast, her skill is low ($z < z^*$), she does not exert effort and becomes an analyst, receiving wage w. Applying equation (3.12), the definition of the cutoff z^* in (3.13), and the labor market clearing condition ($Q = S = z^*$), the welfare function is simplified as follows:

$$\begin{aligned} U_h &= \int_{z^*}^{1} \left(\alpha\rho\sigma_v^2\beta^2(z^*)z + \frac{\left(\alpha\rho\sigma_v^2\beta^2(z^*)\frac{\phi z^*}{2} - w\right)^2}{2\gamma} - \kappa \right) dz + \int_0^{z^*} w dz \\ &= \alpha\rho\sigma_v^2\beta^2(z^*)\frac{(1-z^*)^2}{2} + w. \end{aligned} \tag{5.1}$$

In this expression, w in the second term represents the level of the expected utility that an agent always receives regardless of her occupation, and she obtains additional utility in the first term if she becomes a firm.

In contrast, if the economy is in the decentralized equilibrium (indicated by the subscript "d"), the expected surplus is expressed as

$$\begin{aligned} U_d &= \int_{z_d^*}^{1} \left(\alpha\rho\sigma_v^2\beta^2(z_d^*)z - \kappa \right) dz \\ &= \alpha\rho\sigma_v^2\beta^2(z_d^*)\frac{(1-z_d^*)^2}{2}. \end{aligned}$$

Intuition follows from (5.1), while the wage income is not available due to the absence of the labor market activity, and the cutoff is characterized by z_d^*.

Finally, the noise trader obtains the following expected utility:

$$\begin{aligned} U_n &= E[(v-p)u] \\ &= -\lambda\sigma_u^2, \end{aligned} \tag{5.2}$$

as she executes her order of size u and incurs the price impact λ for that trading. The functional form of U_n is common across the two equilibria, while the price impact λ differs across the equilibria.

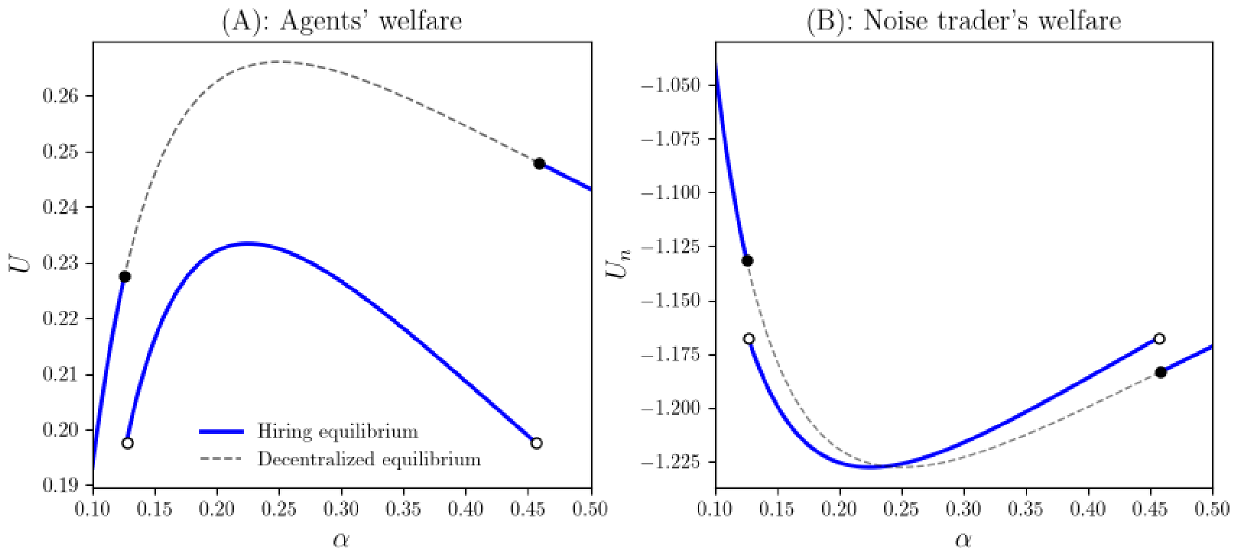

Figure 5: Welfare

Note: This figure plots an agent's expected welfare and that of the noise trader. The parameter setting is the same as that used in Figure 1.

5.1 Welfare Implications of Organized Information Production

A natural question is whether the emergence of organized information production through the analyst labor market improves welfare. Figure 5 plots the agents' expected utilities in two different regimes, U_h and U_d, and the utility of the noise trader, U_n, against the technology level α, suggesting that organized information production does not necessarily improve welfare.

Compared with a hypothetical decentralized equilibrium at the same level of technology, the hiring equilibrium reduces the *ex-ante* welfare of agents. The intuition is that hiring enables managers to aggregate information held by low-skilled agents, who would otherwise be inactive in the absence of the labor market. By utilizing these agents, the organized information production increases the total amount of information incorporated into financial market, as captured by $\tau_d(z^*) < \tau(z^*)$ at the same level of the cutoff. While this improves price informativeness, it also intensifies competition among informed trading firms and compresses informational rents. Individual firms fail to internalize this externality when deciding how many analysts to hire, leading to excessive information production from the perspective of the aggregate agents' welfare.

The analyst wage in U_h partially offsets this welfare loss by allowing low-skilled informed agents, who would otherwise be unable to profitably trade on their own, to monetize their information. However, this benefit is necessarily limited in equilibrium. By construction, the wage must remain sufficiently low to make hiring profitable for managers. Consequently, the labor market cannot generates sufficient compensation to offset the reduction in informational rents caused by intensified competition. As a result, organized information production lowers the *ex-ante* welfare of agents even though it emerges as a competitive equilibrium.

This result contrasts with the traditional view that specialization and the division of labor improve welfare by increasing productive efficiency (e.g., Becker and Murphy, 1992). In my model, organizing information production indeed allows firms to utilize information that would otherwise remain idle. However, unlike standard production environments, financial information is valuable only when it generates informational rents in competitive markets. As organized information production expands, these rents are competed away through a market-wide information externality. My model therefore suggests that, unlike in conventional theories of organizational design, the emergence of information-production organizations does not necessarily imply a welfare improvement for the agents who participate in them.

5.2 Policy Implications

Broadly speaking, governments may influence the organization of financial information production and the equilibrium outcome through two distinct channels. The first is technology policy, including subsidies for AI adoption, public support for financial technologies, or regulations that restrict the use of AI in investment activities. Such policies primarily affect the technological frontier by shifting the economy's information-processing capability. The second is organizational policy, which directly influences firms' incentives to organize information production through hiring, employment regulations, or other policies affecting the cost and scale of research organizations.

Our numerical analysis suggests that neither type of intervention is likely to generate monotonic welfare redistribution. In particular, agents' and the noise trader's welfare exhibits a non-monotonic relationship with technological progress, reflecting the endogenous response of information production and market competition. Consequently, policies that promote financial information technologies primarily redistribute welfare between informed traders and noise traders, rather than generating uniform welfare gains. While welfare redistribution is a standard implication of zero-sum trading, our model predicts that this redistribution is itself non-monotonic: as information technology improves, the welfare of both agents and the noise trader can change non-monotonically.

Organizational policies may have a qualitatively different effect. As discussed above, organized information production intensifies the information externality by encouraging firms to aggregate otherwise dispersed information. Policies that reduce firms' incentives to organize information production may therefore mitigate this externality. In particular, when multiple equilibria exist, interventions that weaken organized information production may induce a transition from the hiring equilibrium to the decentralized equilibrium. Our numerical results suggest that such a transition can improve the welfare of both informed agents and liquidity traders under some parameter values, indicating that organizational interventions may sometimes achieve welfare improvements that are difficult to obtain through technology policy alone.

6 Conclusion

This paper studies how technological progress reshapes the organization of financial information production through the interaction between labor markets and financial markets. I show that technological progress can generate multiple and opposing organizational changes: it initially decentralizes information production by expanding individual capabilities, then induces concentration by intensifying competition among informed firms, and eventually leads to a second form of decentralization when organized information production becomes unprofitable. The key mechanism is that competitive information externalities affect not only the profitability of information production but also the organizational structure through which information is produced. These results highlight that the relationship between technology and organizational structure is more complex than the conventional view that technological progress simply empowers individuals.

Several extensions remain for future research. The model abstracts from several features of real-world financial markets, including heterogeneous technologies across firms and strategic technology adoption. Future work could examine how firms endogenously choose whether and how to adopt new information technologies, and how such decisions interact with market competition and organizational design. More broadly, understanding how general-purpose technologies reshape the boundary between individuals and organizations remains an important question for the future of information-intensive industries.

References

Admati, A. R. (1985). A noisy rational expectations equilibrium for multi-asset securities markets. *Econometrica 53*(3), 629–657.

Aquilina, M., E. Budish, and P. O'Neill (2021). Quantifying the high-frequency trading "arms race". *The Quarterly Journal of Economics 137*, 493–564.

Autor, D., D. Dorn, L. Katz, C. Patterson, and J. Van Reenen (2020). The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics 135*, 645–709.

Becker, G. S. and K. M. Murphy (1992). The division of labor, coordination costs, and knowledge. *Quarterly Journal of Economics 107*(4), 1137–1160.

Bloomberg (2026). New hedge funds are using ai bots to rival industry giants. Accessed: 2026-07-07.

Brynjolfsson, E., D. Li, and L. R. Raymond (2023). Generative ai at work. *Science 381*(6654), 170–172.

Business Insider (2025). Hedge funds' growing divide. Accessed: 2026-07-07.

Chang, A., X. Dong, X. Margin, and C. Zhou (2026). Ai democratization and trading inequality. *Journal of Accounting Research*, 1287–1331.

Farboodi, M. and L. Veldkamp (2020). Long-run growth of financial data technology. *American Economic Review 110*(8), 2485–2523.

Forbes (2026). How ai is leveling the playing field between retail and institutional traders. Accessed: 2026-07-07.

Fukui, M. (2018). Asset quality cycles. *Journal of Monetary Economics*, 97–108.

Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of Political Economy 108*(5), 874–904.

Garicano, L. and E. Rossi-Hansberg (2006). Organization and inequality in a knowledge economy. *Quarterly Journal of Economics 121*(4), 1383–1435.

Glode, V., R. C. Green, and R. Lowery (2012). Financial expertise as an arms race. *Journal of Finance 67*(5), 1723–1759.

Grossman, S. J. and J. E. Stiglitz (1980). On the impossibility of informationally efficient markets. *American Economic Review 70*(3), 393–408.

Holden, C. W. and A. Subrahmanyam (1992). Long-lived private information and imperfect competition. *The Journal of Finance 47*(1), 247–270.

Korinek, A. (2023). Language models and cognitive automation for economic research. *Journal of Economic Literature 61*(4), 1282–1338.

Kurlat, P. (2013). Lemons markets and the transmission of aggregate shocks. *American Economic Review 103*(4), 1463–1489.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 1315–1335.

Sherwin, R. (1981). The economics of superstars. *The American Economic Review 71*, 845–858.

Stein, J. C. (2002). Information production and capital allocation: Decentralized versus hierarchical firms. *Journal of Finance 57*(5), 1891–1921.

Stein, J. C. (2009). Presidential address: Sophisticated investors and market efficiency. *Journal of Finance 64*(4), 1517–1548.

Stiglitz, J. E. and A. Weiss (1981). Credit rationing in markets with imperfect information. *American Economic Review 71*(3), 393–410.

Veldkamp, L. L. (2006). Media frenzies in markets for financial information. *American Economic Review 96*(3), 577–601.

Verrecchia, R. E. (1982). Information acquisition in a noisy rational expectations economy. *Econometrica 50*(6), 1415–1430.

Vissing-Jørgensen, A. (2002). Towards an explanation of household portfolio choice heterogeneity: Nonfinancial income and participation constraints. *NBER Working Paper* (8884).

Vives, X. (2008). *Information and Learning in Markets: The Impact of Market Microstructure*. Princeton, NJ: Princeton University Press.

Wilson, C. (1980). The nature of equilibrium in markets with adverse selection. *Bell Journal of Economics 11*(1), 108–130.

Zandt, T. V. (1999). Real-time decentralized information processing as a model of organizations with boundedly rational agents. *Review of Economic Studies 66*(3), 633–658.

Appendix

A Proof of Propositions 5 and 6

If the hiring equilibrium exists, the cutoff z^* is the solution to $G(z,\alpha)=0$ with

$$G(z,\alpha)=R(z,\alpha)z\left(1-\frac{\phi}{2}\right)+\gamma\frac{z\left(1-\frac{z}{2}\right)}{(1-z)^2}-\kappa, \tag{A.1}$$

where $R(z,\alpha)=\alpha\rho\sigma_v^2\beta(\tau(z,\alpha))$, and β is the unique solution to $B(\beta,\tau(z))=0$ with

$$B(\beta,\tau(z))=\rho\beta(\beta^2\tau^2(z)\sigma_v^2+\sigma_u^2)-\sigma_u^2. \tag{A.2}$$

Since $\frac{\partial B}{\partial\beta}>0$ and $\frac{\partial B}{\partial z}=-\alpha(1-\phi)z\frac{\partial B}{\partial\tau}<0$, the implicit function theorem implies $\frac{d\beta}{dz}>0$. Therefore, $\frac{\partial G}{\partial z}>0$, $G(0,\alpha)=-\kappa<0$, and $G(1,\alpha)=\infty$ hold, suggesting that $G=0$ has a unique solution $z^*\in(0,1)$. Given z^*, the labor market clearing condition yields the following wage:

$$w(z^*)=\frac{\phi}{2}R(z^*,\alpha)z^*-\gamma\frac{z^*}{1-z^*}.$$

If $w(z^*)>0$, then z^* is the hiring equilibrium cutoff. Otherwise, there is no hiring equilibrium.

The decentralized equilibrium is defined by the cutoff z_d^* as the solution to $G_d(z,\alpha)=0$, where

$$G_d(z,\alpha)=R(z,\alpha)z-\kappa. \tag{A.3}$$

The definition of R is the same as above, while $\beta(\tau_d(z))$ involves $B(\beta,\tau_d(z))=0$ with $\tau_d(z)=\frac{\alpha(1-z^2)}{2}$. Due to the same arguments as the hiring equilibrium, $\frac{\partial G}{\partial z}>0$. It also holds that $G_d(0,\alpha)=-\kappa<0$ and $G_d(1,\alpha)=R(1,\alpha)$. Note that $\tau_d(1)=0$ and $\beta(\tau_d(1))=\frac{1}{\rho}$, leading to $G_d(1,\alpha)=\frac{\alpha\sigma_v^2}{\rho}-\kappa$. The solution to $G_d(z,\alpha)=0$ exists in $z\in(0,1)$ if, and only if, $G_d(1,\alpha)>0$, that is, $\alpha>\frac{\rho\kappa}{\sigma_v^2}$. Since Assumption 2 ensures that $\alpha_0>\frac{\rho\kappa}{\sigma_v^2}$, the decentralized equilibrium always exists.

B Proof of Lemma 7 and Propositions 8–10

As $\frac{\partial R}{\partial z} > 0$ holds, it suffices to evaluate the following partial derivative:

$$\frac{\partial R}{\partial \alpha} = \rho\sigma_v^2\left(\beta(\tau(z^*,\alpha)) + 2\alpha\frac{\partial\beta(\tau(z^*,\alpha))}{\partial\alpha}\right). \tag{B.1}$$

Note that $\frac{\partial\beta(\tau(z^*,\alpha))}{\partial\alpha} = \frac{\partial\tau}{\partial\alpha}\frac{\partial\beta(\tau)}{\partial\tau}$, where equation (3.16) implies that

$$\frac{\partial\tau}{\partial\alpha} = \frac{1}{2}(1-(1-\phi)z^{*2}) = \frac{\tau}{\alpha}. \tag{B.2}$$

Also, from equation (A.2), $\tau^2 = \frac{\sigma_u^2}{\sigma_v^2}\frac{1-\rho\beta}{\rho\beta^3}$. By taking the partial derivative of both sides with respect to α,

$$2\tau\frac{\partial\tau}{\partial\alpha} = -\frac{\sigma_u^2}{\sigma_v^2}\frac{(3-2\rho\beta)}{\rho\beta^4}\frac{\partial\beta}{\partial\alpha}.$$

Applying (B.2) and rearranging for $\partial\beta/\partial\alpha$ yields

$$\frac{\partial\beta}{\partial\alpha} = -\frac{2\beta(1-\rho\beta)}{\alpha(3-2\rho\beta)} < 0.$$

Then, the partial derivative (B.1) becomes

$$\frac{\partial R}{\partial\alpha} = \frac{\rho\sigma_v^2\beta(\tau(z^*,\alpha))}{3-2\rho\beta(\tau(z^*,\alpha))}\left(2\rho\beta(\tau(z^*,\alpha))-1\right).$$

Since $\beta < 1/\rho$ holds in the equilibrium, the numerator is positive, leading to

$$\frac{\partial R}{\partial\alpha} \gtreqless 0 \Leftrightarrow \beta(\tau(z^*,\alpha)) \gtreqless \frac{1}{2\rho}.$$

As β is the solution to $B = 0$ with $\frac{\partial B}{\partial\beta} > 0$, the last inequality is re-written as follows:

$$\begin{aligned}\beta > \frac{1}{2\rho} &\Leftrightarrow B(1/2\rho,\alpha) = \frac{1}{2}\left(\frac{1}{4\rho^2}\tau(z^*,\alpha)^2\sigma_v^2 + \sigma_u^2\right) - \sigma_u^2 < 0,\\ &\Leftrightarrow \tau(z^*,\alpha) < 2\rho\frac{\sigma_u}{\sigma_v}\\ &\Leftrightarrow z^* > z_0(\alpha),\end{aligned} \tag{B.3}$$

where

$$z_0(\alpha) \equiv \begin{cases} z_0^+(\alpha) = \sqrt{\frac{1-\frac{1}{\alpha}\frac{4\rho\sigma_u}{\sigma_v}}{1-\phi}} & \text{if } \alpha \geq \frac{4\rho\sigma_u}{\sigma_v}, \\ 0 & \text{if } \alpha < \frac{4\rho\sigma_u}{\sigma_v}. \end{cases}$$

By the implicit function theorem, z^* is decreasing in α if $z^* > z_0(\alpha)$, and vice versa.

Firstly, if $\alpha > \frac{4\rho\sigma_u}{\phi\sigma_v}$, then $\alpha > \frac{4\rho\sigma_u}{\sigma_v}$ holds, and $z_0^+ > 1$ implies that condition (B.3) is always violated. In this case, $\beta < 1/2\rho$ always holds, and thus z^* is increasing in α. Secondly, if $\alpha \leq \frac{4\rho\sigma_u}{\sigma_v}$, then $z_0(\alpha) = 0$, suggesting that $z^* > z_0(\alpha)$ always holds. Thus, in this case, z^* is decreasing in α. Thirdly, consider the case with $\frac{4\rho\sigma_u}{\sigma_v} < \alpha \leq \frac{4\rho\sigma_u}{\phi\sigma_v}$. In this case, $z_0(\alpha) = z_0^+(\alpha) < 1$. Note that z_0^+ is monotonically increasing in α, while z^* is increasing in α if and only if $z^* < z_0$. Thus, consider $\alpha = \alpha^*$ such that $z^*(\alpha^*) = z_0^+(\alpha^*)$. If such α^* exists, then it is a uniquely determined, as $\alpha < \alpha^*$ implies $z^* > z_0$ and vice versa, i.e., z^* and z_0 crosse at most once. As$z^*(\alpha^*) = z_0^+(\alpha^*)$ is equivalent to $2\rho\beta = 1$, equation (A.1) implies that α^* solves $g(\alpha^*) = 0$ with

$$g(\alpha) \equiv \alpha\sigma_v^2 z_0^+(\alpha)\left(1 - \frac{\phi}{2}\right) + 4\rho\gamma\frac{z_0^+(\alpha)\left(1 - \frac{z_0^+(\alpha)}{2}\right)}{\left(1 - z_0^+(\alpha)\right)^2} - 4\rho\kappa. \tag{B.4}$$

g is monotonically increasing in α^*. Also, as $\alpha^* \to \frac{4\rho\sigma_u}{\phi\sigma_v}$, it holds that $z_0^+ \to 1$, and thus $g \to \infty$; as $\alpha^* \to \frac{4\rho\sigma_u}{\sigma_v}$, it holds that $z_0^+ \to 0$, and thus $g \to -4\rho\kappa < 0$. Therefore, within the interval $\frac{4\rho\sigma_u}{\sigma_v} < \alpha \leq \frac{4\rho\sigma_u}{\phi\sigma_v}$, $g(\alpha) = 0$ always has a unique solution, which we define α^*. Following the arguments so far,

$$\frac{dz^*}{d\alpha} \lesseqgtr 0 \Leftrightarrow \alpha \lesseqgtr \alpha^*.$$

By the implicit function theorem, $\frac{\partial R}{\partial \alpha} \gtreqless 0 \Leftrightarrow \frac{dz^*}{d\alpha} \lesseqgtr 0$. This concludes the proof of Lemma 7 and Proposition 8. As $\mu_A = \frac{\phi z^*}{2}$, the above inequality implies point (ii) of Proposition 9.

From the equilibrium condition (A.1),

$$\begin{aligned} R(z^*, \alpha) &= \alpha\rho\sigma_v^2\beta^2(\tau(z^*, \alpha))z^* \\ &= \frac{2\kappa - \gamma\left(\frac{1}{(1-z^*)^2} - 1\right)}{2 - \phi} \\ &\equiv \bar{R}(z^*), \end{aligned}$$

where α does not directly appear in the last expression. From equation (3.13), wage $w(z^*)$ satisfies

$$\bar{R}(z^*) + \frac{\left(\bar{R}(z^*)\frac{\phi_l}{2} - w\right)^2}{2\gamma} = w + \kappa,$$

so that α affects w only through changes in z^*. It holds that $\frac{dw}{d\alpha} = \frac{dz^*}{d\alpha}\frac{dw}{dz^*}$, and the implicit function theorem implies

$$\frac{\partial w}{\partial z^*} = \frac{d\bar{R}}{dz^*}\frac{\gamma + \left(\bar{R}(z^*)\frac{\phi_l}{2} - w\right)\frac{\phi_l}{2}}{\bar{R}(z^*)\frac{\phi_l}{2} - w + \gamma} < 0,$$

where the last inequality comes from $\frac{d\bar{R}}{dz^*} < 0$. Therefore, $\frac{dw}{d\alpha} \gtreqless 0 \Leftrightarrow \frac{dz^*}{d\alpha} \lesseqgtr 0$, and w exhibits the opposite reaction to changes in α compared to z^*, attesting point (i) of Proposition 9.

As for Proposition 10, equation (3.16) implies that

$$\frac{d\tau}{d\alpha} = \frac{\partial\tau}{\partial\alpha} - \alpha(1-\phi)z^*\frac{dz^*}{d\alpha}.$$

If $\frac{dz^*}{d\alpha} < 0$, then $\frac{d\tau}{d\alpha} > 0$ holds. Suppose that $\frac{dz^*}{d\alpha} > 0$. Note that $z^*(\alpha)$ satisfies $G(z^*(\alpha), \alpha) = 0$. By taking the total derivative of this condition with respect to α and rearranging for $\frac{d\tau}{d\alpha}$, we obtain

$$\begin{aligned}\alpha\rho\sigma_v^2 z^*(\alpha)\left(1-\frac{\phi}{2}\right)\frac{d\tau}{d\alpha}\frac{d\beta}{d\tau} =& -\frac{R(z^*(\alpha),\alpha)}{\alpha}z^*(\alpha)\left(1-\frac{\phi}{2}\right) \\ & - \frac{dz^*}{d\alpha}\frac{\partial}{\partial z^*}\left[R(z^*(\alpha),\alpha)z^*(\alpha)\left(1-\frac{\phi}{2}\right) + \gamma\frac{z^*(\alpha)\left(1-\frac{z^*(\alpha)}{2}\right)}{\left(1-z^*(\alpha)\right)^2}\right] \\ & < 0.\end{aligned}$$

where the last inequality comes from our assumption that $\frac{dz^*}{d\alpha} > 0$. Since $\frac{d\beta}{d\tau} < 0$, it holds that $\frac{d\tau}{d\alpha} > 0$. Moreover, from equation (A.2), $\tau^2 = \frac{\sigma_u^2}{\sigma_v^2}\frac{1-\rho\beta}{\rho\beta^3}$ holds, suggesting that $\frac{d\tau}{d\alpha} \propto \frac{d\beta}{d\alpha}$. Thus, we conclude that $\frac{d\beta}{d\alpha} > 0$. The behavior of λ is computed from (3.5) as

$$\frac{d\lambda}{d\alpha} = \frac{d(\beta\tau)}{d\alpha}\frac{\sigma_u^2 - \beta^2\tau^2\sigma_v^2}{(\beta^2\tau^2\sigma_v^2 + \sigma_u^2)^2}\sigma_v^2.$$

Since $\frac{d(\beta\tau)}{d\alpha} > 0$, it holds that $\frac{d\lambda}{d\alpha} > 0 \Leftrightarrow \beta\tau < \frac{\sigma_u}{\sigma_v}$. Noting that $\beta\tau$ is monotonically increasing in α, Assumption 2 ensures the existence of a unique α^{**} such that $\beta\tau = \frac{\sigma_u}{\sigma_v}$, shaping λ into a single-peaked curve against α.

C Proof of Proposition 11

C.1 The Existence of the Hiring Equilibrium

The hiring equilibrium exists when the solution z^* to $G(z^*, \alpha) = 0$ supports $w > 0$, where we know from Proposition 9 that w takes a hump-shaped curve against α with its peak being $w^* \equiv w(z^*(\alpha^*))$. $G = 0$ implies that

$$\begin{aligned} w(\alpha) &\equiv \frac{\phi}{2-\phi}\left(\kappa - \gamma\frac{z^*(\alpha)\left(1-\frac{z^*(\alpha)}{2}\right)}{\left(1-z^*(\alpha)\right)^2}\right) - \gamma\frac{z^*(\alpha)}{1-z^*(\alpha)} \\ &= \frac{Y(z^*(\alpha))}{(2-\phi)(1-z^*(\alpha))^2}, \end{aligned}$$

where

$$Y(z^*(\alpha)) = \left(\phi\kappa + 2\gamma(1-\frac{\phi}{4})\right)z^{*2}(\alpha) - 2(\phi\kappa+\gamma)z^*(\alpha) + \phi\kappa.$$

At $\alpha = \alpha^*$, it holds that $2\rho\beta = 1$, leading to $z^*(\alpha^*) = z_0^+(\alpha^*) = \sqrt{\frac{1-\frac{1}{\alpha^*}\frac{4\rho\sigma_u}{\sigma_v}}{1-\phi}}$, as defined in Appendix B. Hence, the maximum wage is given by

$$w^* = \left(\phi\kappa + 2\gamma(1-\frac{\phi}{4})\right)z_0^{+2}(\alpha^*) - 2(\phi\kappa+\gamma)z_0^+(\alpha^*) + \phi\kappa.$$

If $w^* > 0$, then the continuity of w^* regarding α implies that there is an interval (α_l, α_h) such that $w(z^*(\alpha)) > 0$.

Consider a quadratic function

$$T(z) = \left(\frac{\phi}{2}\kappa + \gamma(1-\frac{\phi}{4})\right)z^2 - (\phi\kappa+\gamma)z + \frac{\phi}{2}\kappa,$$

so that $w^* = 2T(z_0^+(\alpha^*))$. The discriminant of T is $D = (\phi\kappa+\gamma)^2 - 4\frac{\phi}{2}\kappa\left(\frac{\phi}{2}\kappa + \gamma\left(1-\frac{\phi}{4}\right)\right) = \frac{\gamma\phi^2\kappa}{2} + \gamma^2 > 0$. Also, $T(0) > 0$ and $T(1) = -\gamma\frac{\phi}{4} < 0$, suggesting that $T = 0$ has a unique solution in $z \in (0,1)$, which is

$$z_{w=0} = \frac{\phi\kappa + \gamma - \sqrt{\frac{\gamma\phi^2\kappa}{2} + \gamma^2}}{\phi\kappa + 2\gamma\left(1 - \frac{\phi}{4}\right)}. \tag{C.1}$$

Therefore, $w^* > 0$ holds if and only if $z_0^+(\alpha^*) < z_{w=0}$, which is equivalent to

$$\alpha^* < \alpha_{w=0} \equiv \frac{1}{1-(1-\phi)z_{w=0}^2}\frac{4\rho\sigma_u}{\sigma_v}.$$

Since α^* is a unique solution to $g = 0$ in equation (B.4), $\alpha^* < \alpha_{w=0} \Leftrightarrow g(\alpha_{w=0}) > 0$, that is,

$$\alpha_{w=0}\sigma_v^2\left(1-\frac{\phi}{2}\right)z_{w=0} + 4\rho\gamma\frac{z_{w=0}\left(1-\frac{z_{w=0}}{2}\right)}{(1-z_{w=0})^2} > 4\rho\kappa. \tag{C.2}$$

Moreover, under condition (C.2), the threshold α_j for $w^* > 0$ are implicitly defined by $Y(z^*(\alpha)) = 0$. By applying $w(z^*(\alpha)) = 0$, this condition is re-written as

$$\alpha_j\beta^2(z_{w=0},\alpha_j)(1-z_{w=0}) = \frac{2\gamma}{\phi\rho\sigma_v^2}.$$

C.2 Characterizing the Phase Transition

The decentralized equilibrium solves $G_d(z,\alpha) = 0$ in (A.3), where G_d is monotonically increasing in z. Hence, if

$$G_d(z^*,\alpha) = \alpha\rho\sigma_v^2\beta^2(\tau_d(z^*(\alpha)))z^*(\alpha) - \kappa > 0,$$

then $z^* > z_d^*$ holds. At the threshold α_j, $w = 0$ holds. Thus, the above condition becomes

$$G_{w=0,j} = \alpha_j\rho\sigma_v^2\beta^2(\tau_d(z_{w=0},\alpha_j))z_{w=0} - \kappa > 0.$$

By eliminating α using $G(z,\alpha) = 0$, the above condition becomes equivalent to condition (4.3).

The following figures illustrate the behavior of equilibrium variables when the phase transition from the hiring to the decentralized equilibrium involves an upward jump in the

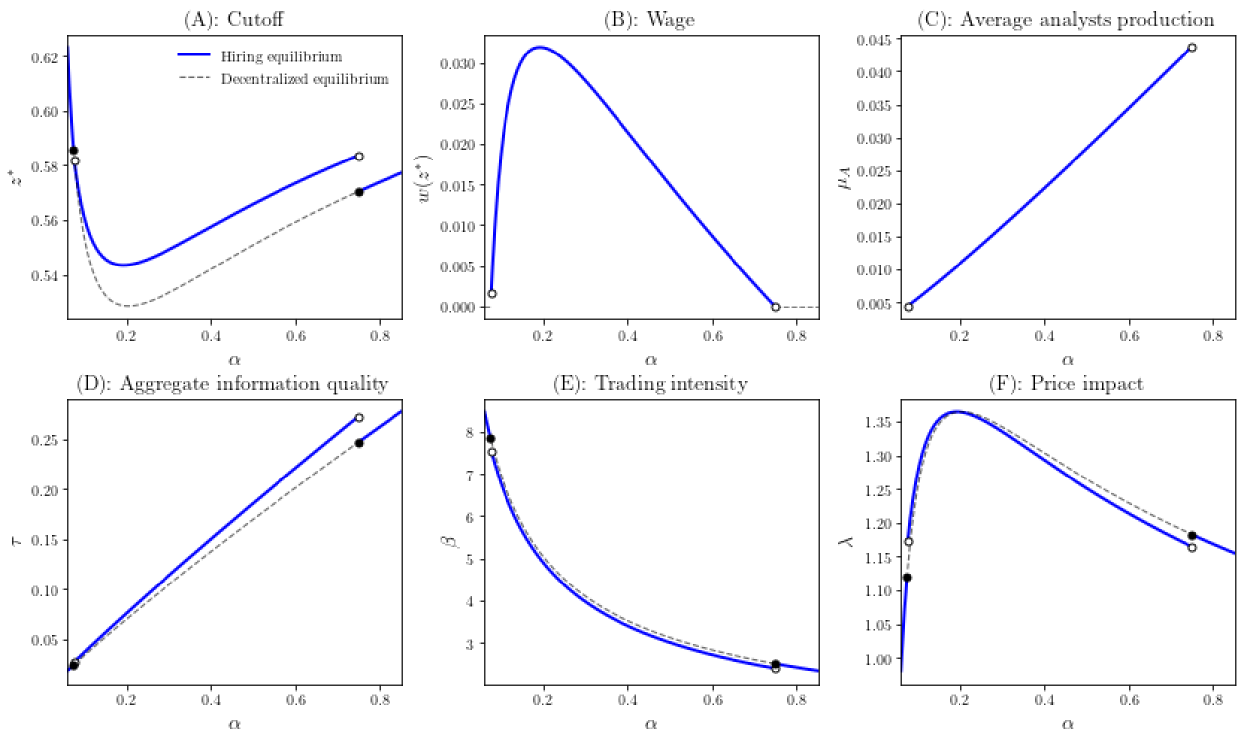

Figure 6: Phase transition

Note: This figure plots the equilibrium variables for the entire range of $\alpha \in [\alpha_0, 1]$. The parameter setting is as follows: $\phi = 0.2, \kappa = 1.8, \gamma = 0.12, \sigma_v^2 = 6.7, \sigma_u^2 = 0.9, \rho = 0.1$.

cutoff z^*.