Variety and Agglomeration in Financial Markets

Ana Babus*

Washington University in St. Louis

July 2026

Abstract

Financial markets are highly differentiated: firms issue many products beyond stocks and bonds. This variety is consequential because the products firms use determine which issuers raise capital, on what terms, and which risks investors absorb. We propose a model of variety and firms' allocation across products. Products repackage firms' cash flows and differ in expected payoff per unit of exposure to aggregate risk and firm-specific risk passed to investors. Aggregate exposure raises the compensation all issuers must offer, whereas firm-specific risk lowers only the issuing firm's payoff. Specialized products coexist with widely adopted ones through trade-offs between these two dimensions.

Keywords: financial product variety; strategic issuance; market structure.

JEL Codes: G12, G23, L13.

*Department of Economics, Washington University in St. Louis, One Brookings Drive, St. Louis, MO 63130. Email: anababus@gmail.com.

1 Introduction

The structure of financial markets shapes how capital and risk are allocated across the economy. Firms raise funds through a wide range of financial products, each packaging cash flows differently. This variety is economically significant. Although a few products, such as common stock and corporate bonds, account for the majority of offerings, many others are adopted by only a small number of firms yet generate substantial funding when used (Babus, Marzani and Moreira, 2026). Which products are used determines which firms raise capital, on what terms, and which risks investors absorb. In conventional product markets, variety is often sustained by heterogeneous consumers with different preferences. Financial markets are different: investors can hold claims issued through multiple products and evaluate them under the same risk–return trade-off. Issuer heterogeneity and customization may explain some variety, but many products are not confined to particular issuer types. What, then, sustains a market structure in which specialized products coexist with widely adopted ones, without product-specific clienteles or issuer-specific customization?

This paper develops a model of financial product variety arising from firms' strategic choices over which products to issue. A product is a contractual technology that maps a firm's uncertain production outcome into the payoff of a financial claim that any investor can trade. Across products, this mapping is characterized by two dimensions: the expected payoff a product delivers per unit of exposure to aggregate risk—the risk shared across firms—and the amount of firm-specific risk it passes through to investors. For example, straight debt promises fixed payments and limits investors' upside, equity passes through residual cash flows, and convertible securities combine debt-like payments with state-contingent participation in equity value. Different contractual mappings can therefore generate trade-offs: a product may deliver more expected payoff per unit of aggregate exposure while also transmitting more firm-specific risk. Products need not be ranked along a single dimension.

The main insight is that the resulting market structure features product variety with uneven agglomeration: some products attract many firms, while others remain specialized. The mechanism comes from how the two dimensions of product differentiation interact with strategic issuance. All else equal, a product that delivers more expected payoff per unit of aggregate exposure is more profitable to issue. At the same time, each issuer faces downward-sloping residual demand: additional issuance requires investors to absorb more risk, and the issuer internalizes the resulting price decline. The two components of that risk operate differently. Firms' uncertain production outcomes are exposed to aggregate risk, but how much of it investors bear depends on the products firms issue: when firms choose products with greater aggregate exposure, more of that risk is transmitted to investors, and the compensation required on every exposed claim rises. This is a common cost, borne by all issuers and shaped by the product choices of all firms. Firm-specific risk instead steepens the issuing firm's own residual demand and lowers her payoff from issuance, without adding to the common risk linking claims across firms: a private cost, borne by the issuing firm alone.

This interaction allows products with different characteristics to coexist. Firms weigh the payoff a product offers against the effect of issuing through it on investors' required compensation for risk. A product that delivers high expected payoff per unit of exposure to aggregate risk attracts firms, but as more firms use it, the overall cost of aggregate risk rises and its advantage shrinks. A product that passes more firm-specific risk to investors instead lowers the payoff of its own issuers by steepening their residual demand. In equilibrium, no product dominates on all relevant margins, so products with different characteristics remain in use side by side. Variety can arise even when all investors can hold all claims and have no product-specific tastes.

The same forces determine how firms distribute across coexisting products. All else equal, firms gravitate toward products with stronger characteristics: higher expected payoff per unit of aggregate exposure and less firm-specific risk passed through to investors. This disciplines which products can remain specialized—that is, adopted by only a few firms. A specialized product cannot dominate a widely adopted one on both dimensions; otherwise, it would attract more firms and cease to be specialized. If it delivers higher expected payoff per unit of aggregate exposure, it must therefore pass through more firm-specific risk. Specialized products reflect trade-offs, not superiority.

The model yields a further implication for variety. When aggregate risk becomes more important, the common component of claim payoffs carries greater weight and the compensation required to absorb it rises. Products with relatively low expected payoff per unit of aggregate exposure then become harder to sustain. Aggregate volatility therefore makes product coexistence more restrictive.

The mechanism makes a distinctive prediction: how many firms adopt a product and how much each raises through it are separate margins, so products chosen by few firms need not raise little when used. This prediction is consistent with evidence in Babus et al. (2026) that specialized products generate proceeds per offering comparable to, and sometimes exceeding, those of widely adopted products.

The theory identifies a source of financial product differentiation that requires neither product-specific investor tastes, nor segmented demand, nor issuer-specific customization. What sustains variety is the incidence of the costs that issuance creates: aggregate exposure imposes a cost common to all issuers, firm-specific risk a cost private to the issuing firm, and products mix the two differently. The skewed market structure of financial products is therefore an equilibrium outcome of common investor demand interacting with firms' strategic product choices.

Related Literature. The paper connects three literatures. The most relevant studies are those showing that market power and market structure in financial markets shape allocations. We build on this view, but study a different margin of market structure: the variety of financial products sustained in equilibrium.

A central theme in the literature on market power in financial markets is that market structure can distort allocations and have real consequences. Models of imperfect competition show how strategic behavior affects prices, allocations, and welfare (Rostek and Weretka, 2015; Malamud and Rostek, 2017). Neuhann and Sockin (2024) show how strategic rationing in illiquid financial markets affects investment and real allocation. We share the view that financial-market structure has allocative consequences. We study a different margin: product variety. Which products coexist determines the claims available to investors and how financing and risk are allocated across firms.

The paper also speaks to work on endogenous market structure and financial innovation. Classic work studies when innovation expands risk-sharing opportunities and whether competition gives agents sufficient incentives to introduce new assets (Allen and Gale, 1991; Carvajal et al., 2012). For a recent survey, see Allen and Barbalau (2024). Recent work studies how market structure interacts with the securities traded. Babus and Hachem (2023) analyze security design when investors choose where to trade, while Rostek and Yoon (2021) study how exchange design and asset characteristics jointly shape risk sharing and efficiency. We contribute by studying which financial products coexist when firms choose among payoff technologies and investors can hold all resulting claims.

This perspective also connects to work on market fragmentation. Babus and Parlatore (2022) study venue choice and Chen and Duffie (2021) analyze fragmentation across trading venues. In these papers, market structure concerns where a given asset trades. In our model, each product defines a market for a different claim, and investors can participate in all product markets. Thus the relevant margin is how issuance is distributed across financial products.

Finally, the paper speaks to the economics literature on product differentiation and product design. Work on multidimensional differentiation shows that coexistence can arise when products cannot be ranked along a single dimension (Irmen and Thisse, 1998; Ansari et al., 1998). Related work studies how firms choose product characteristics or product lines in response to demand (Johnson and Myatt, 2006). In these models, variety is sustained by heterogeneous buyers, product-specific tastes, or locations in a characteristics space. Our paper imports this logic into financial markets, but with a different source of differentiation. Investors have no product-specific tastes, are not segmented, and can hold all claims. Products coexist because they differ in payoff-relevant financial characteristics: the expected payoff they deliver per unit of exposure to aggregate risk and the firm-specific risk they pass through to investors.

This distinction also clarifies the connection to Babus et al. (2026), who document financial product variety and analyze it through the lens of a model with segmented investor demand across products. Here, investors can hold all claims, and coexistence instead arises from firms' product choices and the compensation required to absorb risk.

The rest of the paper proceeds as follows. Section 2 lays out the environment and the equilibrium concept. Section 3 characterizes inverse demand, equilibrium issuance for a fixed allocation of firms across products, and the condition that determines the active set. Section 4 develops the implications for market structure: product variety, specialization, and the allocation of firms across products.

2 Model

2.1 Set-up

Firms. There is a continuum of firms indexed by $\ell \in \Omega \equiv [0, \mu]$. As in Allen and Gale (1988), each firm is endowed with a production plan Y_ℓ, whose realization is uncertain. Uncertainty has two sources: an aggregate factor $\theta \in [0, 1]$, common to all firms, and idiosyncratic shocks, drawn independently across firms conditional on θ. Specifically, conditional on the aggregate factor, firm ℓ's production outcome is drawn from the mixture

$$Y_\ell \sim (1 - \theta)H_1 + \theta H_2, \tag{1}$$

where H_1 and H_2 are arbitrary distributions, common to all firms. For example, θ may summarize the state of aggregate demand, credit conditions, or technological opportunities. When θ is low, firm outcomes are drawn primarily from H_1; when θ is high, they are drawn more heavily from H_2. Thus, the aggregate factor shifts the distribution from which firms' production outcomes are drawn, while idiosyncratic variation within each distribution generates residual differences across firms.

Financial products and claims. A firm raises capital by issuing a financial claim to investors. To do so, she chooses a financial product from a finite set $\mathcal{I}$. A product $i \in \mathcal{I}$ is a contractual technology, represented by a function f_i, that specifies how the payoff promised to investors depends on the realization of the firm's production plan. If firm ℓ uses product

i, the payoff per unit of the resulting claim is

$$W_{i\ell} = f_i(Y_\ell). \tag{2}$$

A product is thus indexed by i alone, whereas a claim is indexed by the pair (i, ℓ): the same product used by different firms generates distinct claims, backed by different production plans.

For any payoff function f_i, the mixture structure in (1) implies that $\mathbb{E}[f_i(Y_\ell) \mid \theta]$ is affine in θ.[1] Hence every claim admits the factor representation

$$W_{i\ell} = x_i + z_i\tilde{\theta} + \varepsilon_{i\ell}, \tag{3}$$

where $\tilde{\theta} \equiv \theta - \mathbb{E}[\theta]$ is the demeaned aggregate factor, with variance σ_θ^2, and $\varepsilon_{i\ell}$ is a firm-specific residual satisfying $\mathbb{E}[\varepsilon_{i\ell} \mid \tilde{\theta}] = 0$ and $\mathbb{V}(\varepsilon_{i\ell}) = \sigma_{\varepsilon_i}^2$. Residuals are uncorrelated across distinct claims. It follows that $\mathrm{Cov}(W_{i\ell}, W_{i'\ell'}) = z_i z_{i'} \sigma_\theta^2$ for $(i, \ell) \neq (i', \ell')$, and $\mathbb{V}(W_{i\ell}) = \sigma_\theta^2 z_i^2 + \sigma_{\varepsilon_i}^2$. Appendix B derives the representation. The unconditional expected payoff is $\mathbb{E}[W_{i\ell}] = x_i$. We assume $x_i > 0$ and $z_i > 0$ for all i: every product delivers a positive expected payoff, and every resulting claim pays more, in expectation, when the aggregate state is high.

Each contractual mapping f_i induces a triple $(x_i, z_i, \sigma_{\varepsilon_i}^2)$ describing the resulting claim's expected payoff, exposure to aggregate risk, and firm-specific residual risk. We summarize these dimensions using two normalized characteristics.

Product i's *systematic productivity*,

$$\chi_i \equiv \frac{x_i}{z_i}, \tag{4}$$

measures the expected payoff the product delivers per unit of exposure to the common factor. Product i's *idiosyncratic share*,

$$\tau_i \equiv \frac{\sigma_{\varepsilon_i}^2}{\sigma_{\varepsilon_i}^2 + \sigma_\theta^2 z_i^2}, \tag{5}$$

is the fraction of the claim's payoff variance accounted for by the firm-specific residual; equivalently, $1 - \tau_i$ is the share explained by the common factor.

Together, χ_i and τ_i make product differentiation two-dimensional. The two characteristics need not be aligned across products: a product with high systematic productivity may also have a high idiosyncratic share. Product choice therefore depends not only on

[1] The only requirement is that $f_i(Y_\ell)$ have finite second moments under H_1 and H_2, so that the residual variance $\sigma_{\varepsilon_i}^2$ is well defined.

systematic productivity but also on the composition of the risk that firms pass to investors.

Product choice and issuance opportunities. Let μ_i denote the mass of firms choosing product i, with $\sum_{i \in \mathcal{I}} \mu_i = \mu$. A product is *active* if it is chosen by a positive mass of firms; the set of active products is $I = \{i \in \mathcal{I} : \mu_i > 0\}$.

Issuance is sparse in the data: most firms issue only rarely (Babus, Marzani and Moreira, 2026). We capture this by assuming that each active product admits a finite number of issuance opportunities, $L_i = L(\mu_i)$, where $L : \mathbb{R}_+ \to \mathbb{N}$ is a weakly increasing step function with $L(0) = 0$. Products chosen by a larger mass of firms thus admit weakly more issuance opportunities, and inactive products admit none. One interpretation is that financial intermediaries can support only finitely many transactions and allocate them across products in response to firms' choices.

We call a product *specialized* when it admits few issuance opportunities and *widely adopted* when it admits many. Our results compare products with fewer and more issuance opportunities without introducing a particular threshold.

Within an active product i, the L_i opportunities are assigned symmetrically among the firms choosing it. Let $\mathcal{L}_i$ denote the finite set of firms assigned issuance opportunities under product i; the assignment rule is formalized in Appendix B. Its relevant implication for the main text is the issuance intensity

$$\iota_i(\mu_i) = \frac{L_i}{\mu_i}, \tag{6}$$

which measures issuance opportunities per unit mass of firms choosing product i.[2] This object is analogous to the meeting intensities in search models of financial markets, such as Duffie, Gârleanu and Pedersen (2005), where pairs of traders randomly meet and receive an opportunity to trade.

A firm $\ell \in \mathcal{L}_i$ that receives an issuance opportunity chooses a quantity $a_{i\ell} \geq 0$ to issue to investors. If each unit of the claim trades at price $p_{i\ell}$, the firm raises proceeds $a_{i\ell} p_{i\ell}$.

Investors. There is a continuum of mean–variance investors of mass one, indexed by n. Investors trade financial claims in competitive markets, taking prices as given. Investor n's holdings in all issued claims are collected in

$$\mathbf{q}^n = \left(q_{i\ell}^n\right)_{i \in I,\ \ell \in \mathcal{L}_i}.$$

Investors are subject to idiosyncratic preference shocks that capture unmodeled determinants of valuation (background holdings, hedging needs, and the like) and that generate

[2] Because firms are nonatomic, $\iota_i(\mu_i)$ should be interpreted as an intensity per unit mass rather than as the probability assigned to any particular firm.

cross-sectional heterogeneity in portfolio choices.

Payoffs and market clearing. We now formalize the payoffs of investors and issuers.

An investor n holding portfolio $\mathbf{q}^n$ expects the following payoff:

$$\mathbb{E}[U^n] = \zeta^n (\mathbf{q}^n)^T \mathbb{E}[\mathbf{W}] - \frac{\gamma}{2} (\mathbf{q}^n)^T \mathcal{V}_{\mathbf{W}} \mathbf{q}^n - \mathbf{p}^T \mathbf{q}^n, \tag{7}$$

where $\mathbf{W} = (W_{i\ell})_{i \in I, \ell \in \mathcal{L}_i}$ is the vector of financial claim payoffs, $\mathcal{V}_{\mathbf{W}}$ is the variance-covariance matrix of the vector of claim payoffs, $\mathbf{p} = (p_{i\ell})_{i \in I, \ell \in \mathcal{L}_i}$ is the vector of prices at which claims trade, $\zeta^n \sim^{i.i.d.} F_\zeta$ is a preference shock with mean $\bar{\zeta} < \infty$, and $\gamma > 0$ controls risk aversion. We assume $\bar{\zeta} > 1$: on average, investors value expected payoffs more than issuers do, so gains from trade exist between firms and the investors who fund them.

The price $p_{i\ell}$ is the price at which the market for claim $W_{i\ell}$ clears, given the amount $a_{i\ell}$ that firm ℓ issues. Thus,

$$\int q_{i\ell}^n(\mathbf{p}) dn = a_{i\ell} \quad \forall i \in I \text{ and } \ell \in \mathcal{L}_i. \tag{8}$$

A firm $\ell \in \mathcal{L}_i$ issuing $a_{i\ell}$ units of claim $W_{i\ell}$ at price $p_{i\ell}$ expects a payoff

$$\mathbb{E}[\Pi_{i\ell}] = a_{i\ell}\big(p_{i\ell} - \mathbb{E}[W_{i\ell}]\big), \tag{9}$$

so $\mathbb{E}[W_{i\ell}]$ is the marginal cost of issuing the claim.[3]

At the product-choice stage, firms value a product by scaling the payoff conditional on receiving an issuance opportunity by the issuance intensity. The relevant object is therefore

$$V_{i\ell} = \iota_i(\mu_i) \cdot \mathbb{E}[\Pi_{i\ell}] = \frac{L_i}{\mu_i} \, a_{i\ell}\big(p_{i\ell} - \mathbb{E}[W_{i\ell}]\big). \tag{10}$$

Because firms are nonatomic, $V_{i\ell}$ is interpreted as a payoff per unit mass of firms choosing product i. This is the object that enters the product-choice problem. A firm that chooses an inactive product obtains payoff zero.

Timing. The timing is as follows. First, firms choose which financial product to issue, generating the mass allocation $\mathcal{M} = \{\mu_i\}_{i \in \mathcal{I}}$. Second, issuance opportunities are assigned within each active product, determining the sets $\{\mathcal{L}_i\}_{i \in \mathcal{I}}$ of firms that receive them. Third, each selected firm $\ell \in \mathcal{L}_i$ chooses an issuance quantity $a_{i\ell} \geq 0$, investors choose portfolios, and prices clear markets.

[3]We assume throughout that selected firms have deep pockets and always make the promised payment $W_{i\ell}$ to investors. Claims are thus default-free, as in Carvajal et al. (2012). The model isolates how products transform underlying production risk into investor payoffs.

2.2 Equilibrium Concept

We define the equilibrium concept in three steps. First, investors choose portfolios, and market clearing generates the inverse demand faced by each issuing firm. Second, taking as given the distribution of firms across products, each firm that receives an issuance opportunity chooses how much to issue. Finally, firms choose among products, determining both the set of active products I and the mass allocation $\{\mu_i\}_{i\in\mathcal{I}}$.

Definition 1 (Equilibrium)**.** *An* equilibrium *consists of a portfolio allocation* $\mathbf{q}^n$ *for each investor* n*, a vector of quantities* $\mathbf{a} = (a_{i\ell})_{i\in\mathcal{I},\ell\in\mathcal{L}_i}$ *supplied by firms* $\ell \in \mathcal{L}_i$*, and a distribution of firms across products* $\mathcal{M} = \{\mu_i\}_{i\in\mathcal{I}}$*, such that the following conditions hold:*

(i) $\mathbf{q}^n$ *solves each investor* n*'s problem*

$$\max_{\mathbf{q}^n} \left\{ \zeta^n (\mathbf{q}^n)^T \mathbb{E}[\mathbf{W}] - \frac{\gamma}{2} (\mathbf{q}^n)^T \mathcal{V}_{\mathbf{W}} \mathbf{q}^n - (\mathbf{p})^T \mathbf{q}^n \right\}.$$

where each price $p_{i\ell}$ *satisfies market clearing, Eq. (8), for each* ℓ*;*

(ii) $a_{i\ell}$ *maximizes each issuer* ℓ*'s expected profit from the issuance*

$$\max_{a\geq 0} a \left[p_{i\ell}(a, \mathbf{a}_{-(i,\ell)}) - \mathbb{E}(W_{i\ell}) \right],$$

where $p_{i\ell}(a, \mathbf{a}_{-(i,\ell)})$ *is the inverse residual demand for claim* $W_{i\ell}$ *implied by Eq. (8);*

(iii) $\mathcal{M} = \{\mu_i\}_{i\in\mathcal{I}}$ *is such that no issuer* ℓ *of product* i *benefits from deviating and switching to another product* i'*. Equivalently, for every* $i, i' \in \mathcal{I}$ *and every* ℓ *choosing* i *then*

$$V_{i\ell}(\mathcal{M}) \;\geq\; V_{i'\ell}(\mathcal{M}).$$

This set-up separates firms' strategic issuance decisions from their product-choice decisions. At the issuance stage, a firm is large and therefore internalizes its price impact. In choosing the quantity $a_{i\ell}$ to issue, firm ℓ takes into account how the price of claim $W_{i\ell}$ depends on its own issuance and on the quantities issued by the other firms in the economy.

At the product-choice stage, however, an individual deviation from product i to product i' changes neither the mass $\mu_{i'}$ of firms choosing product i' nor its issuance intensity $\iota_{i'}$. The deviation payoff is therefore evaluated at the prevailing equilibrium value density in product i'. This separation keeps the product-choice problem tractable while allowing firms that receive issuance opportunities to internalize their price impact in financial markets.

3 Equilibrium

We now characterize equilibrium outcomes, beginning with inverse demand and then turning to issuance and product choice.

3.1 Inverse demand

We begin with investors' portfolio problem, taking as given the set of active products, the assignment of issuance opportunities, and the quantities issued by firms. Our object of interest is the inverse demand system as a function of the quantities supplied.

Investor n's optimal portfolio satisfies the first-order condition

$$\mathbf{q}^n = \frac{1}{\gamma}\mathcal{V}_{\mathbf{W}}^{-1}\left[\zeta^n \mathbb{E}[\mathbf{W}] - \mathbf{p}\right]. \tag{11}$$

This expression reflects the standard mean–variance trade-off. Investors value claims with higher expected payoffs, but require compensation for the risk created by the portfolio of claims supplied to the market.

Proposition 1 (Inverse demand)**.** *Given an assignment of issuance opportunities $\{\mathcal{L}_i\}_{i\in I}$ and quantities issued by firms, let $A_j \equiv \sum_{\ell'\in\mathcal{L}_j} a_{j\ell'}$ denote total issuance in product j. Then the inverse demand for claim $W_{i\ell}$ is*

$$p_{i\ell}(a_{i\ell}, \mathbf{a}_{-(i,\ell)}) = \bar{\zeta}\mathbb{E}[W_{i\ell}] - \gamma\left[\sigma_{\varepsilon_i}^2 a_{i\ell} + z_i^2\sigma_\theta^2 A_i + z_i\sigma_\theta^2\sum_{j\neq i} z_j A_j\right]. \tag{12}$$

Proposition 1 decomposes a claim's price into a valuation component and a risk discount. The valuation component, $\bar{\zeta}\mathbb{E}[W_{i\ell}]$, reflects investors' average valuation of the claim's expected payoff. The risk discount is the compensation investors require to absorb the risks embedded in the claims supplied to the market. To interpret it, let $W_I \equiv \mathbf{a}'\mathbf{W}$ denote the payoff on the aggregate portfolio of claims. Since $\mathrm{Cov}(\tilde{\theta}, W_I) = \sigma_\theta^2\sum_{j\in I} z_j A_j$, the risk discount in (12) can be written as

$$\gamma\left[\sigma_{\varepsilon_i}^2 a_{i\ell} + z_i\ \mathrm{Cov}(\tilde{\theta}, W_I)\right]. \tag{13}$$

This representation reveals the asymmetry between the two components of risk. Firm-specific risk affects only the price of the corresponding claim, whereas aggregate risk affects the prices of all claims exposed to the common factor. A firm's own issuance lowers the price of her claim through both components. It requires investors to absorb more of the claim's firm-specific risk and, at the same time, increases investors' exposure to the common factor. The first effect is governed by $\sigma_{\varepsilon_i}^2$, while the second depends on the claim's loading z_i: the larger z_i, the more strongly the firm's issuance contributes to the aggregate-risk discount. Issuance by other firms affects the price only through the common factor. Issuance through one product therefore raises the compensation investors require on claims issued through other products. The discount generated by firm-specific risk is private to the issuing firm; the discount generated by aggregate risk is common across issuers and is larger for claims with greater exposure to the common factor.

This decomposition clarifies the role of specialized products. A specialized product is chosen by fewer issuers, so a firm issuing through it faces less price pressure from others using the same product. Pressure from the rest of the market instead depends on the claim's loading on the common factor and operates whether the product is chosen by many firms or by few. Specialization therefore limits only within-product price pressure.

3.2 Equilibrium issuance for a fixed allocation

Next, we characterize the issuance decision of a firm that receives an opportunity to issue, taking as given the set of active products. A firm that issues a claim faces a downward-sloping inverse demand curve and therefore internalizes the effect of its issuance on the price of the claim. Thus, firm ℓ, issuing claim $W_{i\ell}$ through product i, chooses $a_{i\ell}$ to satisfy

$$p_{i\ell}(a_{i\ell}, \mathbf{a}_{-(i,\ell)}) - \mathbb{E}[W_{i\ell}] + \frac{\partial p_{i\ell}(a_{i\ell}, \mathbf{a}_{-(i,\ell)})}{\partial a_{i\ell}} a_{i\ell} = 0, \tag{14}$$

where $p_{i\ell}(a_{i\ell}, \mathbf{a}_{-(i,\ell)})$ is the inverse demand in Equation (12). The term $\mathbb{E}[W_{i\ell}]$ is the expected payment promised to investors and is therefore the marginal cost of issuing the claim. The last term in (14) captures the firm's price impact: additional issuance lowers the price that the firm can obtain for her claim. Combining this condition with inverse demand yields the following characterization.

Proposition 2 (Fixed-allocation equilibrium)**.** *Fix the set of active products I and let $\mathcal{X}(I) \equiv \gamma(\bar{\zeta}-1)^{-1} \operatorname{Cov}(\tilde{\theta}, W_I)$. Then there exists a unique interior equilibrium. The amount issued by each firm $\ell \in \mathcal{L}_i$ satisfies the fixed-point condition*

$$a^*_{i\ell} = \frac{\bar{\zeta}-1}{\gamma} \frac{1-\tau_i}{\sigma^2_\theta z_i(1+\tau_i)} \left[\chi_i - \mathcal{X}(I)\right], \tag{15}$$

with expected payoff

$$\mathbb{E}[\Pi^*_{i\ell}] = \frac{(\bar{\zeta}-1)^2}{\gamma\sigma^2_\theta} \frac{1-\tau_i}{(1+\tau_i)^2} \left[\chi_i - \mathcal{X}(I)\right]^2, \tag{16}$$

and $\mathcal{X}(I)$ solves to

$$\mathcal{X}(I) = \frac{\sum_{j\in I} \chi_j L_j \frac{1-\tau_j}{1+\tau_j}}{1 + \sum_{j\in I} L_j \frac{1-\tau_j}{1+\tau_j}}. \tag{17}$$

Proposition 2 shows how product characteristics and market structure jointly shape issuance. The payoff from issuing product i depends first on its own characteristics: systematic productivity, χ_i, and idiosyncratic share, τ_i. A higher χ_i raises the claim's expected payoff per unit of exposure to the common factor. All else equal, this increases issuance and the payoff from issuance. The idiosyncratic share, τ_i, works in the opposite direction. A higher τ_i makes the price of the claim more sensitive to the firm's own issuance. Issuance therefore becomes more costly at the margin, so both the equilibrium quantity and the payoff from issuance fall, all else equal.

The payoff from issuing claim $W_{i\ell}$ also depends on the rest of the market through $\mathcal{X}(I)$. The index $\mathcal{X}(I)$ gains the interpretation of a *market productivity index* as it aggregates the systematic productivities of active products, with product j receiving weight $L_j(1-\tau_j)/(1+\tau_j)$. Products issued by more firms receive more weight; products with higher idiosyncratic shares receive less. A widely issued product with a low idiosyncratic share therefore has the largest effect on the environment all issuers face, while a specialized product—especially one passing through more idiosyncratic risk—has the smallest. This comparison is what firms weigh when choosing products, which we turn to next.

3.3 Product choice and equilibrium variety

The previous section characterized issuance for a fixed set of active products. We now turn to product choice. Because firms choose products before issuance opportunities are assigned, they balance the payoff conditional on issuance against the rate at which issuance opportunities arise. Product differentiation in (χ_i, τ_i) determines the payoff from issuing, while issuance intensity disciplines the allocation of firms across products. Together, these forces determine the distribution of firms across products and discipline the set of products that can remain active. The following proposition formalizes the resulting restriction.

Proposition 3 (Active-set equilibria)**.** *In any equilibrium that supports a set $I \subseteq \mathcal{I}$ of active products, the following condition must hold:*

$$\chi_i > \mathcal{X}(I) \quad \textit{for all } i \in I. \tag{18}$$

The condition has a simple interpretation. Every active product must clear the market productivity index generated by the set of active products. If a product's systematic productivity lies below the index, the quantity in Equation (19) is negative. Given the nonnegativity constraint on issuance, a firm then chooses to issue nothing. The two-product case makes this logic transparent, allows us to construct equilibria in which both products coexist, and identifies the threshold beyond which this coexistence fails.

Take two products, $\mathcal{I} = i, i'$, with $\chi_i \geq \chi_{i'} > 0$, and let $L(m) = L_{i'}$ for $0 < m < \mu/2$ and $L(m) = L_i$ for $m \geq \mu/2$, where $L_i \geq L_{i'} \geq 1$. Consider the distribution

$$\mu_i^* = \mu \, \frac{L_i \mathbb{E}(\Pi_i)}{L_i \mathbb{E}(\Pi_i) + L_{i'} \mathbb{E}(\Pi_{i'})}, \qquad \mu_{i'}^* = \mu - \mu_i^*,$$

where $\mathbb{E}(\Pi_i)$ and $\mathbb{E}(\Pi_{i'})$ are the continuation-equilibrium payoffs. This distribution makes firms indifferent between the two products. When $L_i \mathbb{E}(\Pi_i) > L_{i'} \mathbb{E}(\Pi_{i'})$, it implies $\mu_i^* > \mu/2 > \mu_{i'}^*$, consistent with the step function assigning L_i opportunities to product i and $L_{i'}$ opportunities to product i'.

Under this allocation, both products clear the index $\mathcal{X}(i, i')$, and hence coexist in equilibrium, provided the productivity gap is not too large:

$$\chi_{i'} > \frac{L_i \frac{1-\tau_i}{1+\tau_i}}{1 + L_i \frac{1-\tau_i}{1+\tau_i}} \chi_i. \tag{19}$$

Since $\chi_i \geq \chi_{i'}$, product i clears the index whenever product i' does, so (19) is the binding condition. If it fails, the constructed equilibrium cannot sustain both products: any positive mass on product i' leaves $\chi_{i'} \leq \mathcal{X}(i, i')$, so its issuers earn zero and switch to product i.

The example also clarifies the source of multiplicity. Because firms form a continuum, a product chosen by zero mass receives no issuance opportunities, and a unilateral deviation by a single firm does not change this. This does not mean that every set of products can be supported. A product with sufficiently low systematic productivity cannot remain active, so the cutoff condition restricts which products can coexist. By contrast, the continuation equilibrium is unique for any fixed set of active products (Proposition 2). Multiplicity therefore arises only from firms' product choices.

4 Variety and Agglomeration

The equilibrium characterization now delivers the two features of market structure described at the outset. Specialized products can be used alongside widely adopted ones, even though all claims are priced by the same investors. And firms are distributed unevenly across the products in use: a few attract most issuers, while many are used by only a handful. We take them in turn.

4.1 Variety and specialization

We begin with variety. The equilibrium market structure is jointly determined: each product must clear the market productivity index, but each product also helps determine the index faced by the others. Product j enters $\mathcal{X}(I)$ with weight $L_j(1-\tau_j)/(1+\tau_j)$. A higher χ_j raises the index, whereas a higher τ_j lowers product j's weight in it. The following corollary formalizes this effect.

Corollary 1 (Idiosyncratic share and the market productivity index)**.** *Given a set of active products I, for every product $j \in I$,*

$$\frac{\partial \mathcal{X}(I)}{\partial \tau_j} < 0.$$

The corollary shows how firm-specific risk can support greater variety. A higher idiosyncratic share reduces product j's weight in the market productivity index and lowers the benchmark that other products must clear. Thus, when a product with a higher τ_j is issued, products with more modest systematic productivity can be issued alongside it. In the two-product case, this is the force captured by Equation (19): the cutoff faced by product i' falls as τ_i rises.

Aggregate volatility operates through the same mechanism in reverse. Since $\tau_j = \sigma^2_{\varepsilon_j}/(\sigma^2_{\varepsilon_j} + \sigma^2_\theta z_j^2)$, a rise in σ^2_θ lowers the idiosyncratic share of every active product and raises the market productivity index, holding each χ_j fixed. A more volatile common factor therefore raises the benchmark that products must clear and makes coexistence harder to sustain, especially for products whose systematic productivity lies close to the index.

Coexistence does not yet determine which products are specialized. Product characteristics determine the payoff from issuing, while issuance intensity determines how firms distribute across the products in use. A specialized product is chosen by fewer firms and admits weakly fewer issuance opportunities. Because those opportunities are spread over a smaller mass, however, its issuance intensity can exceed that of a widely adopted product. When it does, firms may accept a lower payoff conditional on issuing in exchange for more frequent issuance opportunities. This trade-off restricts the characteristics of specialized products.

Corollary 2 (Specialization and product characteristics)**.** *Consider an equilibrium with two active products s and w, where product s is specialized and product w is widely adopted, so that $\mu_s < \mu_w$. Suppose also that*

$$\frac{L(\mu_s)}{\mu_s} \geq \frac{L(\mu_w)}{\mu_w}.$$

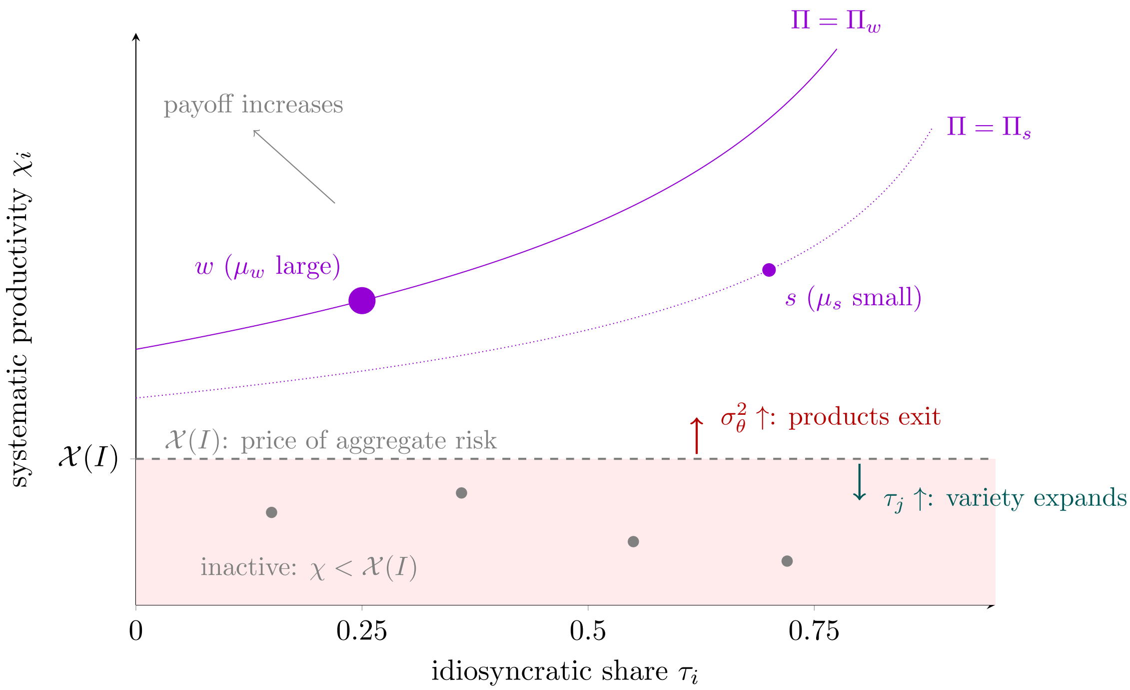

Figure 1: **Market structure in the characteristics plane.** The figure shows an equilibrium with two active products: a widely adopted product w and a specialized product s. Each product is represented by its characteristics (τ_i, χ_i), and dot size reflects the mass of firms choosing it. The dashed line is the market productivity index $\mathcal{X}(I)$; only products above it can be active. The curves are level sets of the per-issuer payoff.

If the specialized product has higher systematic productivity, $\chi_s > \chi_w$, then it must also have higher idiosyncratic share, $\tau_s > \tau_w$.

The corollary shows that a specialized product cannot dominate a widely adopted product on both dimensions. If product s offered both higher systematic productivity and lower idiosyncratic share, it would deliver a higher payoff from issuing while also offering at least as many issuance opportunities per firm. In that case, firms would not choose product w. A specialized product can therefore be stronger on one dimension only if that advantage is offset on the other.

Figure 1 summarizes this trade-off in the space of product characteristics. Each product is a point (τ_i, χ_i). The dashed line is the market productivity index; the gray points below it are products that cannot be sustained.[4] The line is flat: a product's own idiosyncratic share moves it horizontally, not across the boundary. The payoff curves slope upward, because a

[4]Comparing a product against the index generated by the active set is the correct entry test: $\chi_j > \mathcal{X}(I)$ if and only if $\chi_j > \mathcal{X}(I \setminus \{j\})$. Thus, a product fails against the index including it exactly when it fails against the index excluding it.

higher idiosyncratic share makes the firm's own price impact steeper, so keeping the payoff fixed requires higher systematic productivity.

Since all active products give firms the same overall value, and product s has weakly higher issuance intensity, it lies on a weakly lower payoff curve: firms accept a lower payoff conditional on issuing because they issue more often. Because the payoff curves slope upward, a specialized product can have higher systematic productivity only if it also has a higher idiosyncratic share. This is exactly the restriction in Corollary 2.

4.2 Agglomeration

We now turn to the second feature of market structure: the uneven distribution of firms across the products in use. Products with stronger payoff characteristics attract more firms. As firms cluster on a product, its issuance opportunities are spread over a larger mass, lowering issuance intensity until firms are indifferent across products. Thus, for any active set I and any $i, j \in I$,

$$\frac{L_i}{\mu_i}\frac{1-\tau_i}{(1+\tau_i)^2}\left[\chi_i - \mathcal{X}(I)\right]^2 = \frac{L_j}{\mu_j}\frac{1-\tau_j}{(1+\tau_j)^2}\left[\chi_j - \mathcal{X}(I)\right]^2. \tag{20}$$

Equation (20) determines how firms distribute across products. A product whose systematic productivity lies farther above the benchmark, or that passes through less firm-specific risk, delivers a higher payoff conditional on issuing and therefore attracts more firms. The larger mass lowers its issuance intensity until firms again obtain the same payoff from choosing each product. Differences in product characteristics are therefore reflected in how heavily products are represented in equilibrium.

Proposition 4 (Firm allocation across products)**.** *Consider two active products i and j with the same number of issuance opportunities, $L_i = L_j = L$. Holding idiosyncratic share fixed, the product with higher systematic productivity attracts more firms. Holding systematic productivity fixed, the product that passes through less firm-specific risk attracts more firms.*

To see the result, Equation (20) implies

$$\frac{\mu_i}{\mu_j} = \frac{\frac{1-\tau_i}{(1+\tau_i)^2}\left[\chi_i - \mathcal{X}(I)\right]^2}{\frac{1-\tau_j}{(1+\tau_j)^2}\left[\chi_j - \mathcal{X}(I)\right]^2}.$$

Holding $\tau_i = \tau_j$, the product with higher χ_i lies farther above the benchmark and attracts more firms. Holding $\chi_i = \chi_j$, the product with lower τ_i delivers a higher payoff from issuing and attracts more firms.

The same two characteristics that govern coexistence therefore also determine how heavily each product is represented in equilibrium. When they pull in opposite directions, the allocation depends on their relative strength. A product with high systematic productivity may remain specialized if it also passes through substantial firm-specific risk, as in Corollary 2. Adoption and issuance when a product is used are therefore distinct margins: a product chosen by few firms need not generate little issuance when used.

5 Conclusion

This paper shows how financial product variety and uneven adoption can arise from strategic issuance under common investor demand. Products differ in the expected payoff they deliver per unit of exposure to aggregate risk and in the firm-specific risk they pass through to investors. Because investors price all claims jointly, aggregate exposure creates a common cost across issuers, whereas firm-specific risk creates a private cost for the issuing firm. Products that combine these costs differently can therefore coexist even without segmented demand, product-specific tastes, or issuer-specific customization.

The analysis also highlights a broader implication of strategic issuance. Investors can trade all claims, but the claims available to them remain determined by firms' product and issuance choices. Because each product bundles aggregate and firm-specific risk in fixed proportions, the equilibrium set of claims need not span all payoff directions, so markets are generically incomplete. Market structure therefore shapes not only which firms obtain funding, but also the risks investors can trade.

References

Allen, F., Barbalau, A., 2024. Security design: A review. Journal of Financial Intermediation 60, 101113. doi:10.1016/j.jfi.2024.101113.

Allen, F., Gale, D., 1988. Optimal security design. The Review of Financial Studies 1, 229–263. doi:10.1093/rfs/1.3.229.

Allen, F., Gale, D., 1991. Arbitrage, short sales, and financial innovation. Econometrica 59, 1041–1068. doi:10.2307/2938173.

Ansari, A., Economides, N., Steckel, J., 1998. The max-min-min principle of product differentiation. Journal of Regional Science 38, 207–230.

Babus, A., Hachem, K., 2023. Markets for financial innovation. Journal of Economic Theory 208, 105615. doi:10.1016/j.jet.2022.105615.

Babus, A., Marzani, M., Moreira, S., 2026. The Rise of Specialized Financial Products. Technical Report. Working Paper. Manuscript, version cited in text.

Babus, A., Parlatore, C., 2022. Strategic fragmented markets. Journal of Financial Economics 145, 876–908. doi:10.1016/j.jfineco.2021.08.022.

Carvajal, A., Rostek, M., Weretka, M., 2012. Competition in financial innovation. Econometrica 80, 1895–1936. doi:10.3982/ECTA9837.

Chen, D., Duffie, D., 2021. Market fragmentation. American Economic Review 111, 2247–2274. doi:10.1257/aer.20200829.

Duffie, D., Gârleanu, N., Pedersen, L.H., 2005. Over-the-counter markets. Econometrica 73, 1815–1847. doi:10.1111/j.1468-0262.2005.00639.x.

Irmen, A., Thisse, J.F., 1998. Competition in multi-characteristics spaces: Hotelling was almost right. Journal of Economic Theory 78, 76–102.

Johnson, J.P., Myatt, D.P., 2006. On the simple economics of advertising, marketing, and product design. American Economic Review 96, 756–784. doi:10.1257/aer.96.3.756.

Malamud, S., Rostek, M., 2017. Decentralized exchange. American Economic Review 107, 3320–3362. doi:10.1257/aer.20140759.

Neuhann, D., Sockin, M., 2024. Financial market concentration and misallocation. Journal of Financial Economics 159, 103875. doi:10.1016/j.jfineco.2024.103875.

Rostek, M., Weretka, M., 2015. Dynamic thin markets. The Review of Financial Studies 28, 2946–2992.

Rostek, M., Yoon, J.H., 2021. Exchange design and efficiency. Econometrica 89, 2887–2928. doi:10.3982/ECTA16537.

Appendix

A Proofs

Proof of Proposition 1. From the factor representation $W_{i\ell} = x_i + z_i\tilde{\theta} + \varepsilon_{i\ell}$ (Appendix B), with $\mathbb{E}[\varepsilon_{i\ell} \mid \tilde{\theta}] = 0$ and residuals uncorrelated across distinct claims, the payoff covariance is $\mathrm{Cov}(W_{i\ell}, W_{i'\ell'}) = z_i z_{i'}\sigma_\theta^2 + \mathbf{1}\{(i,\ell) = (i',\ell')\}\sigma_{\varepsilon_i}^2$; stacking claims, $\mathcal{V}_\mathbf{W} = \sigma_\theta^2 \mathbf{z}\mathbf{z}' + \bar{\Sigma}$, where $\mathbf{z}$ repeats z_i for each $\ell \in \mathcal{L}_i$ and $\bar{\Sigma} = \mathrm{diag}(\sigma_{\varepsilon_i}^2)$.

Investor n's first-order condition gives $\mathbf{q}^n = \gamma^{-1}\mathcal{V}_\mathbf{W}^{-1}[\zeta^n \mathbb{E}\mathbf{W} - \mathbf{p}]$. Integrating over investors with $\int \zeta^n\, dn = \bar{\zeta}$ and imposing market clearing $\int \mathbf{q}^n\, dn = \mathbf{a}$ yields $\mathbf{p} = \bar{\zeta}\,\mathbb{E}\mathbf{W} - \gamma\,\mathcal{V}_\mathbf{W}\mathbf{a}$. With $\mathcal{V}_\mathbf{W}\mathbf{a} = \sigma_\theta^2\mathbf{z}(\mathbf{z}'\mathbf{a}) + \bar{\Sigma}\mathbf{a}$, the (i,ℓ) component is $\sigma_\theta^2 z_i \sum_{j\in I} z_j A_j + \sigma_{\varepsilon_i}^2 a_{i\ell}$, where $A_j = \sum_{\ell'\in\mathcal{L}_j} a_{j\ell'}$. Splitting the sum into its own- and cross-product parts gives (12). $\square$

Remark 1 (Auxiliary facts)**.** *Write $\omega_j \equiv L_j\frac{1-\tau_j}{1+\tau_j}$, $S_{-j} \equiv \sum_{k\neq j}\omega_k\chi_k$, $\Omega_{-j} \equiv \sum_{k\neq j}\omega_k$, so that $\mathcal{X}(I) = (\omega_j\chi_j + S_{-j})/(1+\omega_j+\Omega_{-j})$. Cancelling $\omega_j\chi_j$ from both sides of $\chi_j > \mathcal{X}(I)$ gives*

$$\chi_j > \mathcal{X}(I) \iff \chi_j > \mathcal{X}(I \setminus \{j\}) \equiv \frac{S_{-j}}{1+\Omega_{-j}}.$$

A product therefore fails against the index including it exactly when it fails against the index excluding it.

Proof of Proposition 2. Each issuer's objective $a\,[p_{i\ell}(a, \mathbf{a}_{-(i,\ell)}) - \mathbb{E}W_{i\ell}]$ is strictly concave in a, with second derivative $-2\gamma(\sigma_{\varepsilon_i}^2 + z_i^2\sigma_\theta^2) < 0$. Stacking the Kuhn–Tucker conditions, a profile $\mathbf{a} \geq 0$ is an equilibrium if and only if it solves the complementarity problem $\mathbf{b} - \gamma M\mathbf{a} \leq 0$, $\mathbf{a}'(\mathbf{b} - \gamma M\mathbf{a}) = 0$, with $\mathbf{b} = (\bar{\zeta}-1)\mathbb{E}\mathbf{W}$ and $M = \mathcal{V}_\mathbf{W} + \bar{\Sigma} + \sigma_\theta^2\,\mathrm{diag}(z_i^2)$. Since $\mathcal{V}_\mathbf{W}$ is positive semidefinite and the remaining terms are positive diagonal, M is symmetric positive definite, and the solution is unique. We now compute it and give the condition under which it is interior.

Combining (14) with (12), using $a_{i\ell}\,\partial p_{i\ell}/\partial a_{i\ell} = -\gamma(\sigma_{\varepsilon_i}^2 + z_i^2\sigma_\theta^2)a_{i\ell}$ and $\mathbb{E}W_{i\ell} = x_i = z_i\chi_i$,

$$(\bar{\zeta}-1)z_i\chi_i = \gamma\Big[z_i\sigma_\theta^2 \textstyle\sum_{j\in I} z_j A_j + (2\sigma_{\varepsilon_i}^2 + z_i^2\sigma_\theta^2)a_{i\ell}\Big]. \tag{A.1}$$

Within-product symmetry gives $a_{i\ell} = a_i^*$, so $A_i = L_i a_i^*$ and $\sum_{j\in I} z_j A_j = \sum_{j\in I} L_j z_j a_j^*$. The right-hand side depends on the profile only through the scalar $K \equiv \frac{\gamma\sigma_\theta^2}{\bar{\zeta}-1}\sum_{j\in I} L_j z_j a_j^*$. Solving (A.1) for a_i^*,

$$a_i^* = \frac{\bar{\zeta}-1}{\gamma}\,\frac{z_i}{2\sigma_{\varepsilon_i}^2 + z_i^2\sigma_\theta^2}\big[\chi_i - K\big]. \tag{A.2}$$

Multiplying by $L_i z_i$, summing, and using $\sigma_\theta^2 z_i^2/(2\sigma_{\varepsilon_i}^2 + z_i^2\sigma_\theta^2) = (1-\tau_i)/(1+\tau_i)$ gives

$K = \sum_i L_i \frac{1-\tau_i}{1+\tau_i}\chi_i - K\sum_i L_i \frac{1-\tau_i}{1+\tau_i}$, hence

$$K = \frac{\sum_{i\in I} L_i \frac{1-\tau_i}{1+\tau_i}\chi_i}{1+\sum_{i\in I} L_i \frac{1-\tau_i}{1+\tau_i}} = \mathcal{X}(I).$$

Substituting $K = \mathcal{X}(I)$ into (A.2) and using $z_i/(2\sigma_{\varepsilon_i}^2 + z_i^2\sigma_\theta^2) = [\sigma_\theta^2 z_i]^{-1}(1-\tau_i)/(1+\tau_i)$ gives (15). This profile is the (unique) solution of the complementarity problem, and it is interior, if and only if $\chi_i > \mathcal{X}(I)$ for every $i \in I$. Finally, the first-order condition gives $p_{i\ell}^* - \mathbb{E}W_{i\ell} = \gamma(\sigma_{\varepsilon_i}^2 + z_i^2\sigma_\theta^2)a_{i\ell}^*$; with $\sigma_{\varepsilon_i}^2 + z_i^2\sigma_\theta^2 = \sigma_\theta^2 z_i^2/(1-\tau_i)$ and (15),

$$p_{i\ell}^* - \mathbb{E}[W_{i\ell}] = (\bar\zeta - 1)\,\frac{z_i}{1+\tau_i}\big[\chi_i - \mathcal{X}(I)\big], \tag{A.3}$$

and $\mathbb{E}[\Pi_{i\ell}^*] = a_{i\ell}^*(p_{i\ell}^* - \mathbb{E}W_{i\ell})$ yields (16). □

Proof of Proposition 3. Let I be an equilibrium active set and let $I^+ \equiv \{i \in I : a_{i\ell}^* > 0\}$ be the set of products actually supplied. Claims in $I \setminus I^+$ contribute nothing to $\sum_j z_j A_j$, so the argument in the proof of Proposition 2, applied to I^+, gives $K = \mathcal{X}(I^+)$ and $a_{i\ell}^* \propto [\chi_i - \mathcal{X}(I^+)] > 0$ for $i \in I^+$.

First, $I^+ \neq \varnothing$: otherwise $K = 0$, and since $x_i, z_i > 0$ implies $\chi_i > 0$, (A.2) would give $a_{i\ell}^* > 0$ for every $i \in I$, a contradiction. Next, suppose $I^+ \subsetneq I$. Firms choosing $i \in I \setminus I^+$ issue zero and obtain $V_{i\ell} = 0$, whereas (A.3) gives $V_{j\ell} > 0$ for any $j \in I^+$. Such a firm would deviate to j, contradicting condition (iii). Hence $I^+ = I$, so $K = \mathcal{X}(I)$ and $\chi_i > \mathcal{X}(I)$ for all $i \in I$. □

Lemma 1 (Two-product equilibrium)**.** *Let $\mathcal{I} = \{i, i'\}$ with $\chi_i \geq \chi_{i'} > 0$, and let $L(m) = L_{i'}$ for $0 < m < \mu/2$ and $L(m) = L_i$ for $m \geq \mu/2$, with $L_i \geq L_{i'} \geq 1$. Let $\mathbb{E}(\Pi_i)$ and $\mathbb{E}(\Pi_{i'})$ denote the payoffs* (16) *evaluated at $I = \{i, i'\}$. Both products are active in equilibrium if and only if*

$$\chi_{i'} > \frac{\omega_i}{1+\omega_i}\,\chi_i, \qquad \omega_i = L_i\frac{1-\tau_i}{1+\tau_i}, \tag{A.4}$$

and either $L_i = L_{i'}$ or $L_i\mathbb{E}(\Pi_i) > L_{i'}\mathbb{E}(\Pi_{i'})$. The equilibrium allocation is then unique and equal to $\mu_i^* = \mu\, L_i\mathbb{E}(\Pi_i)/[L_i\mathbb{E}(\Pi_i) + L_{i'}\mathbb{E}(\Pi_{i'})]$, $\mu_{i'}^* = \mu - \mu_i^*$.

Proof. By Remark 1 with $j = i'$, condition (A.4) is equivalent to $\chi_{i'} > \mathcal{X}(\{i, i'\})$, and $\chi_i \geq \chi_{i'}$ then gives $\chi_i > \mathcal{X}(\{i, i'\})$. By Proposition 3 this is necessary; by Proposition 2 it delivers a unique continuation equilibrium with $\mathbb{E}(\Pi_i), \mathbb{E}(\Pi_{i'}) > 0$. Condition (iii) requires $\frac{L_i}{\mu_i}\mathbb{E}(\Pi_i) = \frac{L_{i'}}{\mu_{i'}}\mathbb{E}(\Pi_{i'})$ with $\mu_i + \mu_{i'} = \mu$, whose unique solution is $(\mu_i^*, \mu_{i'}^*)$; both sides then equal $[L_i\mathbb{E}(\Pi_i) + L_{i'}\mathbb{E}(\Pi_{i'})]/\mu$. This allocation is consistent with the step function, i.e. $L(\mu_i^*) = L_i$ and $L(\mu_{i'}^*) = L_{i'}$, automatically when $L_i = L_{i'}$, and otherwise exactly when

$\mu_i^* \geq \mu/2 > \mu_{i'}^*$, i.e. when $L_i \mathbb{E}(\Pi_i) > L_{i'} \mathbb{E}(\Pi_{i'})$. If (A.4) fails then $\chi_{i'} \leq \mathcal{X}(\{i, i'\})$, and by Proposition 3 no equilibrium has both products active. $\square$

Proof of Corollary 1. In the notation of Remark 1,

$$\frac{\partial \mathcal{X}(I)}{\partial \omega_j} = \frac{\chi_j(1 + \Omega_{-j}) - S_{-j}}{(1 + \omega_j + \Omega_{-j})^2} = \frac{(1 + \Omega_{-j})\big[\chi_j - \mathcal{X}(I \setminus \{j\})\big]}{(1 + \omega_j + \Omega_{-j})^2} > 0,$$

the sign following from activeness (Proposition 3) and the equivalence in Remark 1. Since $\partial \omega_j / \partial \tau_j = -2L_j/(1 + \tau_j)^2 < 0$, the chain rule gives $\partial \mathcal{X}(I)/\partial \tau_j < 0$.

For aggregate volatility, hold the active set and $\{L_j\}$ fixed. Each $\chi_j = x_j/z_j$ is invariant to σ_θ^2, while $\partial \tau_j / \partial \sigma_\theta^2 = -\sigma_{\varepsilon_j}^2 z_j^2 / (\sigma_{\varepsilon_j}^2 + \sigma_\theta^2 z_j^2)^2 < 0$. Summing the chain rule over $j \in I$, $\partial \mathcal{X}(I)/\partial \sigma_\theta^2 = \sum_j (\partial \mathcal{X}/\partial \omega_j)(\partial \omega_j/\partial \tau_j)(\partial \tau_j / \partial \sigma_\theta^2) > 0$. $\square$

Proof of Corollary 2. Indifference gives $\frac{L(\mu_s)}{\mu_s} \mathbb{E}[\Pi_{s\ell}^*] = \frac{L(\mu_w)}{\mu_w} \mathbb{E}[\Pi_{w\ell}^*]$; with $L(\mu_s)/\mu_s \geq L(\mu_w)/\mu_w > 0$ this gives $\mathbb{E}[\Pi_{s\ell}^*] \leq \mathbb{E}[\Pi_{w\ell}^*]$. Substituting (16), cancelling the common positive constant, and taking square roots (both products active, so $\chi_s - \mathcal{X}(I), \chi_w - \mathcal{X}(I) > 0$),

$$\kappa(\tau_s)\big[\chi_s - \mathcal{X}(I)\big] \leq \kappa(\tau_w)\big[\chi_w - \mathcal{X}(I)\big].$$

If $\chi_s > \chi_w$ then $\chi_s - \mathcal{X}(I) > \chi_w - \mathcal{X}(I) > 0$, so $\kappa(\tau_s) \leq \kappa(\tau_w)\, \frac{\chi_w - \mathcal{X}(I)}{\chi_s - \mathcal{X}(I)} \leq \kappa(\tau_w)$. κ is strictly decreasing, so $\tau_s > \tau_w$. $\square$

Proof of Proposition 4. For active i, j with $L_i = L_j$ (equivalently, μ_i and μ_j lie in the same step of L), indifference (20) and (16) give

$$\frac{\mu_i}{\mu_j} = \frac{g(\tau_i)\, [\chi_i - \mathcal{X}(I)]^2}{g(\tau_j)\, [\chi_j - \mathcal{X}(I)]^2}, \qquad g(\tau) = \frac{1 - \tau}{(1 + \tau)^2},$$

with both brackets positive (Proposition 3). Holding $\tau_i = \tau_j$, $\mu_i > \mu_j \iff \chi_i > \chi_j$. Holding $\chi_i = \chi_j$, $\mu_i > \mu_j \iff g(\tau_i) > g(\tau_j) \iff \tau_i < \tau_j$. $\square$

B Derivations

Assignment of issuance opportunities. Let λ be Lebesgue measure on $\Omega = [0, \mu]$, and let $A_i \subseteq \Omega$ be the set of firms choosing product i, with $\lambda(A_i) = \mu_i$. If $\mu_i = 0$, set $\mathcal{L}_i = \varnothing$. Otherwise assign each of the $L_i = L(\mu_i)$ opportunities independently by a draw from the uniform probability measure $\lambda(\cdot \cap A_i)/\mu_i$ on A_i. Because λ is nonatomic, the L_i draws land on distinct firms almost surely, so $|\mathcal{L}_i| = L_i$; and the expected number of opportunities falling in any measurable $B \subseteq A_i$ is $L_i \lambda(B)/\mu_i$. The induced issuance intensity is therefore

$\iota_i(\mu_i) = L(\mu_i)/\mu_i$, an intensity per unit mass rather than a probability attached to any named firm (each of which is selected with probability zero).

From the mixture structure to the factor representation. Conditional on θ, $Y_\ell \sim (1-\theta)H_1 + \theta H_2$, so for any f_i with finite second moments under H_1, H_2,

$$\mathbb{E}[f_i(Y_\ell) \mid \theta] = (1-\theta)\mathbb{E}_{H_1}[f_i] + \theta\, \mathbb{E}_{H_2}[f_i] = \alpha_i + \beta_i\theta,$$

with $\alpha_i = \mathbb{E}_{H_1}[f_i]$, $\beta_i = \mathbb{E}_{H_2}[f_i] - \mathbb{E}_{H_1}[f_i]$. Defining $\varepsilon_{i\ell} = f_i(Y_\ell) - \mathbb{E}[f_i(Y_\ell) \mid \theta]$ (so $\mathbb{E}[\varepsilon_{i\ell} \mid \theta] = 0$) and centering with $x_i = \alpha_i + \beta_i\mathbb{E}\theta$, $z_i = \beta_i$, $\tilde{\theta} = \theta - \mathbb{E}\theta$, gives $W_{i\ell} = x_i + z_i\tilde{\theta} + \varepsilon_{i\ell}$, which is (3), with $\mathbb{E}W_{i\ell} = x_i$. Orthogonality $\mathrm{Cov}(\tilde{\theta}, \varepsilon_{i\ell}) = \mathbb{E}[\tilde{\theta}\,\mathbb{E}(\varepsilon_{i\ell} \mid \theta)] = 0$ follows from iterated expectations, and $\mathrm{Cov}(\varepsilon_{i\ell}, \varepsilon_{i'\ell'}) = 0$ for distinct claims under conditional independence across firms. The residual is conditionally heteroskedastic under the mixture; only the unconditional variance $\sigma^2_{\varepsilon_i} = \mathbb{V}(\varepsilon_{i\ell})$ enters investors' preferences, so this is immaterial. The maintained sign restriction $z_i > 0$ is the requirement $\mathbb{E}_{H_2}[f_i] > \mathbb{E}_{H_1}[f_i]$: the product pays more, in expectation, in the high aggregate state. Any primitive yielding an affine conditional expectation delivers the same representation.