Information-Concealing Credit Architecture *

Gary Gorton† Ye Li‡ Guillermo Ordoñez§

June 9, 2026

Abstract

Creditors are tempted to examine collateral assets of uncertain value, but costly information acquisition ultimately reduces financing capacity. A pecking order emerges. Debt provides greater financing capacity than equity: unlike equity, creditors own the asset only if the borrower defaults, discouraging costly asset examination. Probabilistic asset ownership can be further diluted by introducing intermediaries between borrower and creditor, leading to a new theory of financial intermediation and credit chains. Our theory rationalizes the seemingly excessive complexity of financial architecture: the optimal chain arises in a decentralized equilibrium and is characterized by a sequence of heterogeneous intermediaries that discourages information production.

*We are very grateful to Marzena Rostek (the editor) and three anonymous referees for their helpful comments to improve the paper. We are also grateful for comments from Philip Bond, Douglas Diamond (discussant), Vincent Glode, Ahmed Guecioueur, Bengt Holmstrom, Doron Levit, Giorgia Piacentino (discussant), Stephan Siegel, Victoria Vanasco, Mark Westerfield, conference and seminar participants at the École Polytechnique Fédérale de Lausanne (EPFL), Finance Theory Group 2026 Spring Meeting, NBER Summer Institute (Corporate Finance), SFS Cavalcade North America 2026, the University of British Columbia, and the University of Washington. The usual waiver of responsibility applies.

†Yale University and NBER. E-mail: gary.gorton@yale.edu

‡University of Washington. E-mail: liye@uw.edu

§University of Pennsylvania and NBER. Email: ordonez@econ.upenn.edu

1 Introduction

The financial system operates through intricate networks of intermediated funding and collateral flows. This complexity of credit intermediation has been widely criticized for its opacity, in part because it obscures systemic vulnerabilities, creates regulatory challenges, and complicates crisis management.[1] Yet, although this seemingly excessive complexity is often framed either as an unavoidable byproduct of risk sharing or as an intentional attempt to hide risk-taking, we argue that it can also benefit the economy by expanding credit capacity.

What is the best way to pledge an asset of uncertain value to obtain credit? When both borrower and lender face that uncertainty, credit capacity is maximized "in the dark": credit based on the expected asset value (an *information-insensitive loan*) exceeds average credit under realized asset values (an *information-sensitive loan*) because symmetric ignorance avoids the cost of information production.[2] The size of an information-insensitive loan is nevertheless limited by an incentive-compatibility constraint: the more credit the lender extends, the stronger the temptation to inspect the collateral rather than lend with "no questions asked."

Throughout, we use the term "asset" to refer both to cash flows generated by non-financial assets that back debt, such as capital equipment and real estate, and to the financial instruments engineered from those cash flows. Prior studies have pointed out that, to deter collateral inspection, an asset is often structured to have low value uncertainty or high examination costs.[3] These approaches, however, are costly and limited. In this paper, we take the asset's characteristics as given and ask how debt contracts and complex intermediation networks can further discourage costly collateral examination. They do so by changing the probability that the lender ultimately acquires the asset: a lender who expects repayment rather than default has little incentive to inspect collateral that she is unlikely to own. We show that debt carries probabilistic asset ownership, and that intermediation based on debt contracts reduces the probabilistic ownership by diffusing it across parties along a credit chain, thereby enhancing the asset's capacity to serve as collateral.

We analyze a setting in which a firm uses an existing asset as collateral to fund a project with high but unpledgeable returns. The goal is to maximize credit against collateral whose value

[1] FRBNY (2020) provides a recent policy discussion, among others. citeAguiar2016ofrwp document a complex map of collateralized credit. Initiatives such as the G20-led push for central clearing and EU's securitization framework reflect a consensus on curbing complexity (e.g., G20 (2018); EU (2022)). Living wills and reforms, such as the UK's ring-fencing regime, highlight crisis management under complex fund flows (e.g., Dallas Fed (2012); PRA (2022)).

[2] There may be other benefits of avoiding information, such as liquidity provision, as in Dang et al. (2017).

[3] Pooling assets into ABS, MBS, and CDO reduces value uncertainty via diversification, and layers of pooling (e.g., CDO-squared) raise inspection costs by increasing opacity. See Gorton and Ordonez (2023) for an overview.

is uncertain to both the firm and investor. Under symmetric ignorance, financing capacity equals the expected asset value. But if financing is large enough, the investor is tempted to examine the collateral, producing information at a cost ultimately borne by the firm if the investor breaks even. The firm must therefore choose between limiting financing to deter examination or accepting costly information production. Relative to equity, debt relaxes this incentive-compatibility constraint: the lender has a weaker incentive to produce information because she owns the collateral only if the firm's project fails and the firm defaults, whereas an equity investor is always exposed to the asset's uncertain value. Debt thus carries *probabilistic asset ownership*, generating a pecking order in which debt supports more information-insensitive financing than equity.[4] We do not present this pecking order itself as a novel contribution, but as a reassuring result. Our contribution is that credit intermediation expands financing capacity by further diluting probabilistic asset ownership.

Consider an intermediary that extends a loan to *an end borrower* and uses that loan as collateral to borrow from *an end lender*. If the borrower defaults and the intermediary receives the collateral, it may still use its own funds to repay the end lender rather than hand over the collateral. As a result, the end lender faces a lower probability of owning the collateral and therefore has a weaker incentive to examine it than if she financed the end borrower directly. The intermediary's incentive to examine is also weaker: the borrower's default does not guarantee that it keeps the collateral, because it may also fail and pass it on to the end lender.

In summary, when the borrower defaults, ownership of the collateral is split probabilistically between the intermediate and end lenders. This ownership dilution weakens both lenders' incentives to examine the collateral relative to a direct lender, relaxing the IC constraints on information-insensitive loans and allowing more credit to flow to the end borrower. Importantly, the intermediary's role is not to contribute new resources that reduce the end lender's exposure to the end borrower's default. In our setting, what drives credit capacity is the financial architecture itself. Consider two firms facing liquidity needs. In this case, credit is maximized when each uses the other as an intermediary rather than borrowing directly from the end lender. Here, credit chains form without requiring additional resources, thereby increasing credit capacity for both firms just because of the chain architecture.

The ownership-dilution logic extends naturally beyond a single intermediary. We embed it in a general model of credit chains with heterogeneous information costs and an arbitrary joint

[4] This differs from Myers and Majluf (1984), whose pecking order responds to a fixed asymmetric-information environment; ours shapes the information environment, as the investor's decision to examine the asset is endogenous.

failure distribution. Each link's credit capacity is shaped by its lender's probabilistic ownership of collateral, which depends on two probabilities: that the lender gets the collateral (everyone upstream fails) and, conditional on getting it, keeps it (she survives conditional on everyone upstream failing). The chain's overall capacity is pinned down by the bottleneck link with the smallest credit capacity. A natural question is which subset of candidate intermediaries, and in what order, maximizes credit capacity. Because the candidate set is finite, a planner's optimum exists, but can it emerge from decentralized choices? We model chain formation as a sequential game in which an initiator—either the end borrower or end lender—invites an intermediary, the two parties bargain bilaterally over the resulting surplus, and the invitee, upon acceptance, becomes the next inviter. We establish that, under both borrower-initiated and lender-initiated formation, a subgame-perfect equilibrium exists in which the chain coincides with the planner's optimum.

Given an existing chain, should it be extended when new agents become available, and if so, how? Enlarging credit capacity requires inserting a new intermediary at the bottleneck. Doing so does not change collateral possession probabilities for preceding intermediaries ("upstream") but strictly reduces them for all subsequent lenders ("downstream"): the probability of getting the collateral falls because joint upstream failure now includes one more agent's failure, and conditional on this more adverse event, the probability of surviving and keeping the collateral also falls. Hence, to evaluate whether a new intermediary should be added, it is enough to compare her incentive to examine the collateral with that of the current bottleneck lender.

When there are many candidates with heterogeneous attributes, such as information costs and conditional survival probabilities, the optimal chain equalizes incentives to examine the collateral across links, avoiding bottlenecks and leaving no slack. The intuition is similar to maximizing a Leontief production function. If intermediaries differ in asset-examination costs, those with higher costs should be placed upstream, closer to the end borrower. If they differ in their probability of survival conditional on getting the collateral, those with lower survival probability (for example, those whose failure is more correlated with the end borrower's) should also be placed upstream. In short, intermediaries that are more likely *to get the collateral* should also be less likely *to keep the collateral*, and vice versa, so that the temptation to examine is equalized along the chain.

A long credit chain, however, is meaningful only when intermediaries are heterogeneous, and their failures are partially correlated. To see why, consider the case in which intermediaries have the same information costs and, conditional on the end borrower's default, independent and identical survival probabilities. The optimal chain features only one intermediary: adding more intermediaries weakens information-production incentives downstream, but the bottleneck remains the first link, whose incentive to examine the collateral is unchanged by the downstream chain.

While our main analysis takes agents' attributes as given, it also has implications for endogenous asset choice. We show that an equilibrium exists in which an agent forgoes productive investments and instead holds less productive but *quality-certain assets* (government bonds, for instance) to operate as a "superior intermediary." Pledging those assets as collateral does not induce information production and is therefore not subject to the IC constraint on financing capacity. At the same time, specializing in quality-certain assets reduces expertise in examining quality-uncertain assets, thereby raising information costs. The first force makes the agent a better borrower, the second a better lender, and together they make the agent a superior intermediary.

Our work provides a rationale for the seemingly excessive complexity of intermediated credit flows. The paper's empirical content is not just that intermediation can raise credit capacity, but that it does so in a sharply structured way. Chains matter more when collateral is information-sensitive and direct lenders are tempted to examine it. Additional intermediaries should be inserted at bottleneck links. Safer or more diversified intermediaries should appear downstream, away from end borrowers, much as in practice, banks finance riskier non-bank lenders that in turn lend to firms.[5] Conversely, when collateral is sufficiently opaque or when intermediaries are homogeneous, long chains add little value. These are distinct predictions of our credit-architecture view.

Our setting naturally applies to funding environments in which collateral values are uncertain, and claimants are tempted to examine them. Non-Treasury collateral has been widely used in credit chains. According to Infante et al. (2018), the collateral pass-through rate—the ratio of rehypothecated collateral to total collateral received—was around 0.9 for non-Treasury assets from 2016 to 2018, implying multiple-link credit chains. Extending the information problem to allow the private value of collateral to depend on the correlation between that collateral and the lender's (potentially opaque) assets, the model also applies when Treasury securities serve as collateral. Finally, we highlight that alternative arrangements for diluting asset ownership and information-acquisition incentives, such as multilateral contracting, may face important frictions, whereas the credit chains we characterize rely only on bilateral debt contracts. Here, it is not the design of complex contracts that enlarges financing capacity, but the architecture of simple contracts.

[5]For household loans, 80% of nonbank lenders' debt is from banks (Jiang, 2023), 40% for BDC loans (Chernenko and Scharfstein (2025) and Haque et al. (2025)) and 25% for REIT debt (Acharya et al., 2025).

Relation with the literature. The problems created by asymmetric information in financing are well known in the literature and date back at least to Hirshleifer (1971) and Rock (1986). Diamond (1985) and Diamond and Verrecchia (1991), for instance, show that when firms have superior information, disclosure can improve borrowing capacity because more credit can be sustained under symmetric information than under asymmetric information. Our paper takes one step back by analyzing an environment in which the information structure is endogenous. We show that more credit can be sustained under symmetric ignorance than under information production that leads to either symmetric or asymmetric information.

This positioning also clarifies the paper's novelty relative to prior studies on information-insensitive debt. We do not present the usefulness of information-insensitive claims as a new result in itself. Our contribution is to show how financing *architecture* changes the object that determines debt capacity: each claimant's probability of owning the collateral is endogenous to the intermediation chain. The question is therefore not only whether debt is preferable to equity, but how a chain of debt contracts should be organized so that informational incentives are diluted where they bind most tightly. Our mechanism is also distinct from cross-pledging, which is about borrowers' skin in the game. Our focus is on lenders' chain-dependent collateral exposure, making the architecture itself part of the mechanism.

Our work also relates to the literature rationalizing the existence of financial intermediaries. Diamond (1984) shows that by holding diversified portfolios, intermediaries reduce delegation costs and improve credit provision. In our model, diversification does not occur *within a bank* but *along a chain*. Moving downstream increases diversification, but that is of little use if the bottleneck is upstream. The optimal chain, therefore, balances an intermediary's probability of getting the asset against its probability of keeping it, or surviving conditional on getting it. In Dang et al. (2017), by contrast, banks improve credit provision by making their assets as opaque as possible. Here, we show that the creation of opaqueness by discouraging costly information production is not only a design of each financial institution but of the whole system.

We show that intermediation chains weaken lenders' incentives to obtain an informational advantage, thereby imposing an information asymmetry. Glode and Opp (2016) and Glode, Opp, and Zhang (2019) take information asymmetry as given and show that intermediation chains mitigate the resulting inefficiency. In their papers, intermediaries (dealers) facilitate asset trading in spot markets, while in our model, intermediaries sign bilateral (debt) contracts. Sannino (2026) shows that, given information asymmetry, intermediaries improve efficiency by allowing sellers to commit to trade.[6] Previous studies also analyze how heterogeneity in liquidity needs or investment opportunities generates interbank credit or insurance networks.[7] Our paper instead shows that heterogeneity in information costs and financial soundness shapes the anatomy of credit chains.

A recent body of literature examines credit chains. Maggio and Tahbaz-Salehi (2014) study how the distribution of collateral along predetermined chains affects funding capacity and systemic stability. In Donaldson and Micheler (2018), credit chains arise to mitigate liquidation losses when banks rely on non-resaleable debt (e.g., repo). In He and Li (2022), the focus is on maturity mismatch and rollover risk, and the intermediation chain shortens debt maturity downstream. Donaldson, Piacentino, and Yu (2022) study the financial stability implications of chains of long-term debt that allow borrowers to dilute existing creditors by issuing new debt. We present here a clear anatomy of credit infrastructure based on the location of intermediaries along the chain.

One form of issuing debt backed by the borrower's assets is a repurchase agreement (repo). Repo requires spot exchange of both cash and collateral, whereas in our model, only cash needs to move from lender to borrower. Issuing equity means raising funds by selling the asset. Under this interpretation, our model supports the superiority of repo over asset sales for raising funds. Our explanation differs from Parlatore (2019), who argues that firms prefer pledging financial assets to selling them when investment returns are unobserved, the asset is illiquid, or investment opportunities are persistent. It also differs from Monnet and Narajabad (2012), who emphasize bilateral trading frictions; in our model, those frictions are absent. The key distinction is between probabilistic asset ownership under debt and full asset ownership under equity or asset sale.

Our paper also contributes to the literature on repo and collateral reuse. Gottardi et al. (2019) rationalizes the use of repos with positive haircuts when investors are risk-averse; our mechanism instead operates under risk neutrality, making it especially relevant when information acquisition is the dominant friction. Infante and Vardoulakis (2021) study the risk of collateral runs arising when cash borrowers may lose their collateral; our mechanism does not rely on this counterparty risk, and the core force remains operative in an extended model with counterparty risk (see Appendix B). Brumm et al. (2023) explain collateral reuse through the scarcity of pledgeable collateral; our mechanism instead operates even when collateral is abundant, because intermediation chains enlarge credit capacity by diluting probabilistic ownership and discouraging costly information

[6] Chains of spot asset exchanges also emerge in models of over-the-counter (OTC) markets, often in the absence of informational frictions (e.g., Viswanathan and Wang (2004); Atkeson et al. (2015); Chang and Zhang (2015); Babus and Kondor (2018); Hugonnier et al. (2019); Hendershott et al. (2020)).

[7] See, for instance, Brusco and Castiglionesi (2007); Babus (2016); Craig and Ma (2022); Farboodi (2023).

production. In Infante and Saravay (2024), the emphasis is on the convenience yield of safe assets, whose reuse as collateral expands effective supply. Our mechanism does not require a special role for safe assets and thus applies naturally to collateral classes such as private-label MBS, corporate loans, and structured securities, where information-production incentives are central.

Finally, our work contributes to the ongoing discussion about the desirability of complex financial systems. While much of this literature emphasizes the negative consequences of complexity, such as Stiglitz (1999) and Caballero and Simsek (2013), we instead highlight a rationale for complexity as a way to improve credit provision in the economy.

The next section lays the contractual foundation for our analysis. It introduces a setting in which costly information production and probabilistic asset ownership under debt financing give rise to a new pecking-order theory. Section 3 shows that intermediation dilutes probabilistic asset ownership and enlarges credit capacity. Section 4 characterizes optimal intermediation chains with heterogeneous intermediaries and shows that the optimum arises in a decentralized equilibrium. In Section 5, we study the rise of specialized intermediaries and their asset choices. The final section concludes. All proofs are in Appendix A.

2 Optimal Financing: A Pecking Order

We analyze a financing problem in which the value of pledgeable assets is uncertain, both for those who demand and for those who supply liquidity. In this section, we focus on one firm obtaining funds from either an equity investor or a creditor. In the subsequent sections, we analyze a richer environment with multiple firms that may act as intermediaries for one another, forming credit chains between the borrower and the end lender. Our environment features debt optimality, but it differs from standard security-design settings—and from the costly-state-verification setting in particular—in that the relevant information concerns the value of pledgeable collateral rather than the outcome of the project seeking financing. Consequently, the key contracting object is not the promised payment but the claimant's probability of ultimately owning the collateral, which links the information problem to financing capacity. This feature allows the analysis to move beyond a debt-versus-equity comparison toward a theory of intermediation chains.

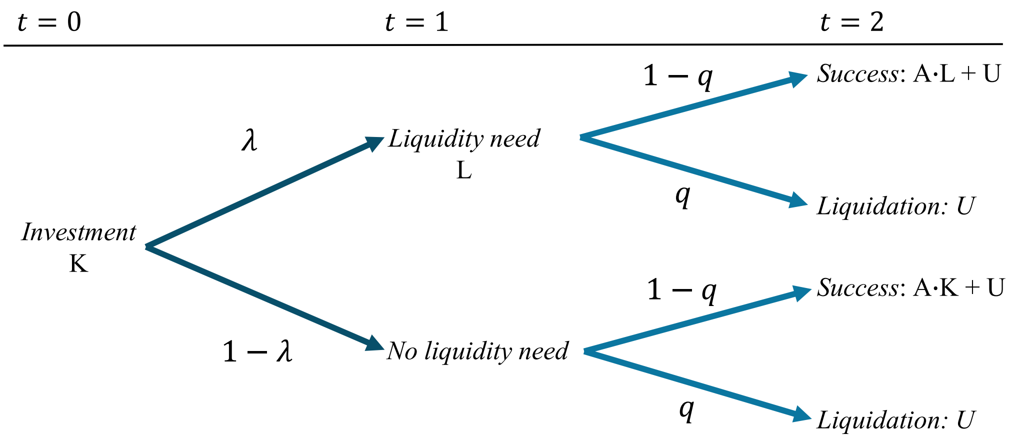

Figure 1: **Project Timeline.**

2.1 The Setup

The economy has three dates, $t = 0$, 1, and 2. There are two agents: a firm and a deep-pocket investor. They are risk-neutral with a zero discount rate.

At $t = 0$, the firm has an endowment K fully invested in a project. The project involves capital equipment, real estate, inventory, and financial assets needed for operations, which represent pledgeable assets with a liquidation value of U at $t = 2$. The project also generates nonpledgeable, stochastic cash flows: with probability q, the project fails, and with probability $(1 - q)$ it generates $A > 1$ per unit of investment. We assume, however, that at $t = 1$ the firm is hit by a liquidity shock with probability λ. In this case, the firm needs to raise new funds to maintain operations, possibly at a reduced scale: if the firm raises $L \in [0, K]$ it operates at a scale that is L/K of the original size, and obtains profits AL if successful. The timeline is illustrated in Figure 1.

We assume that the pledgeable value, U, is uncertain: it is equal to G with probability p and B ($< G$) with probability $1 - p$. Let $\bar{U}$ denote the expected pledgeable value of the project

$$\bar{U} = pG + (1 - p)B. \quad (1)$$

Without the uncertainty in U, the model is rather trivial given the linearity of the setting: the firm would raise up to U in case of a liquidity need to produce AU with probability $1 - q$. We introduce uncertainty in U to meaningfully discuss the examination of U, its implications for whether to obtain funds using equity or debt, and financial architecture (e.g., credit chains).

The information environment is as follows. At $t = 1$ when the liquidity shock hits, the firm does not know whether the liquidation value is G or B. At $t = 1$, the investor does not know the liquidation value either but can learn about it by paying a cost, C.[8] Asset examination cannot be prohibited. Once the investor knows the liquidation value, the firm also learns it. At $t = 2$, the firm knows whether the project succeeds, but the investor does not and cannot learn about the project outcome, either A or 0.[9]

We make several assumptions to highlight the main forces at play. First, the firm has its whole endowment in the project at $t = 0$, so it cannot self-insure against the liquidity shock at $t = 1$. We also rule out external financing at $t = 0$. If it can costlessly pre-fund future liquidity needs, the financing problem at $t = 1$ is resolved. We abstract from this possibility not because building financial slack is not important, but because doing so allows us to focus on the financing problem at $t = 1$ when self-insurance via external financing at $t = 0$ is costly and thus incomplete. Second, the project always has a pledgeable liquidation value of U, independent of the liquidity shock or the project's success. Third, the cash flows from a successful project—AK before the liquidity shock and AL after the liquidity shock—cannot be pledged. The second and third assumptions guarantee that the firm's financing capacity is tied to the value of U, while its incentive to raise financing, L, in the liquidity event is tied to productivity A. Such separation allows a transparent exposition of our mechanism.[10] We also make a couple of parametric assumptions to avoid unnecessary complications and focus on our proposed mechanism.

Assumption 1 *Throughout the analysis, we maintain the following set of parametric conditions:*

- $K \geq G$: *in the liquidity event, financing based on even the highest liquidation value is still insufficient to refinance the project to its original capacity.*
- A *is sufficiently high: the firm wants to maximize refinancing in the liquidity event.*

Starting in the next subsection, we analyze the financing structure in the liquidity event at $t = 1$. In our model, equity represents the investor's *direct ownership* (shares) of U, while debt carries a *probabilistic ownership*—a lender owns U only if the project fails and, as a result, the

[8] The assumption that only the investor has the technology to learn about U is not restrictive but useful for exposition, as discussed by Gorton and Ordonez (2023). The setting can be extended to allow both parties to examine the asset, thereby generating two possible deviations in private information acquisition rather than one. An example of the asymmetry we assume here is the real estate market: the borrower takes out a mortgage from a bank, triggering a credit chain that involves the bank, the repo market, a rehypothecation chain, and money market funds. The borrower does not know the collateral value, nor do the financial professionals (intermediate and end lenders) on the chain.

[9] Note that in the standard costly state verification model, the lender can learn about the project's success or failure at a cost. In contrast, in our model, the investor may learn about the liquidation value U but not the project outcome.

[10] For interpretation, one may regard the liquidation value U as what the investor can repossess, and through the threat of liquidation, what the investor can enforce. Alternatively, one may view A as unobservable and unverifiable, or as inalienable, such as coming from the firm manager's human capital.

firm has no cash flow but only the collateral asset to hand to the lender, which is a probability q event. We show that the probabilistic ownership of collateral assets weakens a lender's incentive to incur the costs of gathering information, so the project's financing capacity is larger under debt than under equity.

Discussion: viability vs. pledgeability. The project's pledgeable value and continuation value differ. The former is U—the uncertain collateral value that tempts costly examination—and the latter is driven by the unpledgeable productivity A. Creditors care about what they can seize or enforce; the borrower and the planner also value the non-pledgeable continuation surplus. The inefficiency is that a high-A project (worth continuing) may not receive sufficient refinancing when hit by liquidity shocks, because the costs of examining U reduce its financing capacity.

If information reveals that a project had negative continuation value (i.e., a low or negative A, which we do not consider in the paper), then suppressing that information would not be desirable. In that case, information has a screening value. Screening firms with good or bad A is not in our setup—we introduce Assumption 1, which states that A is sufficiently high. Our paper is not saying that all information is bad. The point is more localized: once projects are viable (i.e., with a sufficiently high A), costly private information production about collateral quality (U) can reduce financing capacity and induce inefficient liquidation of viable projects.

Therefore, viability is about A, and pledgeability is about U. When it comes to information problems, we focus on those of the latter. An interesting extension may allow a fraction of projects to have negative-NPV (a low A), and this can be detected by a screening technology, then the optimal arrangement should feature a screening of A with the result made public and, in addition, discourage costly private examination of the collateral (U). If the same costly signal reveals both A and U, there is a trade-off: information concealment is valuable only when the financing-capacity benefit for viable (high A) projects dominates the loss from failing to screen out bad (low A) ones.

2.2 Equity Financing

The firm issues equity shares with a liquidation value of U. The number of shares is normalized to one. We consider information-sensitive and information-insensitive equity. In the former, the investor pays the cost C to learn about U. For the investor to break even, the cost of information acquisition must be compensated through a discount in the equilibrium price of equity, implying a financing capacity in expectation less than $\bar{U}$. In the latter, the equity price does not suffer from an information-cost discount, but the financing capacity is limited by an incentive-compatible (IC) condition that prevents the investor from deviating and privately acquiring information about U.

The next lemma summarizes the ex-ante financing capacity for each type of equity, and the ex-ante surplus each generates for the firm. Since we assume the investor breaks even, the firm's surplus is also the social surplus. The proof is in Appendix A.1.

Lemma 1 (Equity financing capacity and ex-ante surplus)

• *Optimal information-sensitive equity contract: The firm raises funds in expectation equal to* $\mathbb{E}(E^s) = \bar{U} - C$ *and generates expected surplus* $(1-q)A(\bar{U} - C) - \bar{U}$ *in the liquidity event.*

• *Optimal information-insensitive equity contract: The firm raises with certainty funds equal to* $E^i = \min\left\{\bar{U}, \Gamma C\right\}$ *, where*

$$\Gamma = \frac{1}{p\left(\frac{G-\bar{U}}{\bar{U}}\right)}, \tag{2}$$

and generates expected surplus $(1-q)AE^i - E^i$ *in the liquidity event.*

When the firm sells shares to raise funds using an information-sensitive contract, the price of equity is contingent on the realization of U minus the information cost, hence in expectation the price is $\bar{U} - C$, where the discount, C, compensates the investor for the information cost. The surplus has two parts. The first part, $(1-q)A(\bar{U}-C)$, is the expected profits reinvesting, where, as previously stated, $1-q$, is the probability that productivity A materializes. The second part, $-\bar{U}$, reflects the fact that the firm pledges the full liquidation value $\bar{U}$, which covers both the investment cost, $\bar{U} - C$, and the investor's information cost, C.

In contrast, when the firm sells shares to obtain funds using an information-insensitive contract, the price of equity is not contingent on the realization of U, and the surplus in this case does not reflect the cost of information acquisition. However, for this contract to be incentive-compatible (IC)—that is, it does not induce the investor's examination—the firm may be restricted to sell only a fraction of the pledgeable value, i.e., $E^i < \bar{U}$ when $\Gamma C < \bar{U}$. Raising funds beyond the IC limit, ΓC, tempts the investor to acquire information because too much is at stake. Notice that Γ captures the properties of U that are inversely related to its *information worthiness*: the investor's incentive to acquire information is stronger when there is a higher probability p of discovering a good asset whose value deviates significantly from the expected value $\frac{G-\bar{U}}{\bar{U}}$. Therefore, financing capacity is greater when Γ is higher (U is less information-worthy) and when the investor's information cost, C, is higher, as both discourage examination of the asset.

2.3 Debt Financing

Now we consider a debt contract that specifies the loan amount, the promised repayment, and a covenant that allows the lender to seize a fraction of the asset (collateral) in the event of default. Again, we consider information-sensitive and information-insensitive debt. In the former, the lender pays the cost C to learn about U, which the borrower must compensate for through the contract. In the latter case, the lending amount is limited by an incentive-compatible (IC) condition that guarantees the lender would not deviate and privately examine the collateral.

The next lemma summarizes the ex-ante credit capacity of each type of debt and the expected social surplus generated by each contract. As in the case of equity financing, the investor (lender) breaks even, so the social surplus fully goes to the firm.

Lemma 2 (Credit capacity and ex-ante surplus)

• *Optimal information-sensitive debt contract: The firm raises funds* $\mathbb{E}(L^s) = \bar{U} - C$ *in expectation and generates an expected social surplus* $(1-q)A(\bar{U}-C) - \bar{U}$ *in the liquidity event.*

• *Optimal information-insensitive debt contract: The firm raises funds* $L^i = \min\left\{\bar{U}, \Gamma\frac{C}{q}\right\}$ *and generates an expected social surplus equal to* $(1-q)AL^i - L^i$ *in the liquidity event.*

Using information-sensitive debt, the firm borrows in expectation of $\bar{U}$ (G if the collateral is found to be good, and B otherwise), net of the information cost C that should be compensated to the lender. The surplus is given by the profits from investing the funds raised, $\bar{U}-C$, net of the fully pledged collateral value, $\bar{U}$, that covers both the investment cost, $\bar{U}$, and the investor's information cost, C. Both the expected financing amount and surplus are equal to the values obtained under an information-sensitive equity contract (see Lemma 1).

In contrast, when the firm borrows using an information-insensitive contract, the lending amount is not contingent on the realization of U, and the surplus does not reflect the cost of information acquisition. However, for this contract to be incentive-compatible, the lending amount cannot exceed $\Gamma\frac{C}{q}$; otherwise, with too much at stake, the investor is tempted to incur costly information costs. If $\Gamma\frac{C}{q} < \bar{U}$, credit capacity is below the pledgeable value.

As in the case of equity, Γ summarizes the attributes of the pledgeable value that discourage information production. This is what we call the *opacity properties of the asset*. Importantly, however, in a debt contract, credit capacity is high when the probability of the project failing is low (i.e., q is low). This is what we call the *repayment properties of the borrower*. When q is low, it is unlikely that the firm will default and the lender will end up with the asset, so the lender is not incentivized to examine it. Importantly, this is not due to the lender's risk aversion and credit risk being priced in equilibrium. In our model, the lender is risk-neutral. The link between default probability and credit capacity emerges from the lender's information choice.

Discussion: debt vs. lines of credit. Alternative financial arrangements can relax the funding constraint at $t = 1$, but they do not, in general, eliminate the informational-acquisition problem. A line of credit provides liquidity insurance, yet if the draw is ultimately backed by the same collateral with uncertain quality, and if the bank that extends the line of credit or the bank's own investors can pay a cost to examine the collateral, the same IC constraint reappears. A line of credit fully neutralizes the friction only under a set of demanding conditions: the commitment must be extended and priced before the information problem arises (i.e., before the bank and its investors are even aware of the uncertainty in collateral quality and the costly technology to examine the collateral) and must remain credible at the draw date, insulated from covenant violations related to collateral revaluation, renegotiation, and the bank's own funding constraints.

The focus on bilateral debt and intermediation chains is deliberate. The paper is not a general security-design exercise over the full contracting space; it studies a class of arrangements that is empirically relevant, debt and credit chains. Debt is central because the mechanism hinges on probabilistic ownership of collateral, which is the key force in our analysis.

2.4 An Informational Theory of the Pecking Order

According to Lemma 1 and 2, information-sensitive debt and equity generate the same financing capacity and social surpluses. Therefore, when information is acquired, our analysis yields a Modigliani-Miller-style result — i.e., debt and equity are equivalent. A meaningful difference between debt and equity emerges when considering information-insensitive contracts. Comparing the financing capacities from information-insensitive equity and debt in Lemma 1 and 2, respectively, we can see that the latter allows for a greater financing capacity: under $q < 1$,

$$\Gamma \frac{C}{q} > \Gamma C. \tag{3}$$

While a lender receives the asset with probability q (when the borrower defaults), an equity investor always owns it. Since producing information is more beneficial when the likelihood of owning the asset increases, equity induces a stronger incentive to produce information than debt and thus has a tighter IC constraint and a smaller financing capacity. If an information-insensitive equity contract is feasible, an information-insensitive debt contract is also feasible. The reverse is not true.

Now we compare the alternatives. When $\Gamma \frac{C}{q} > \Gamma C > \bar{U}$, both information-insensitive debt and equity are feasible and achieve the maximum credit capacity. In the parameter region where $\Gamma \frac{C}{q} > \bar{U} > \Gamma C$, debt generates a greater surplus as it allows borrowing the full pledgeable value $\bar{U}$. Therefore, our model generates a new pecking-order theory—debt is preferred to equity—based on costly information production. In this case, information-insensitive debt also dominates information-sensitive contracts (equity and debt), which only have a financing capacity of $\bar{U} - C$.

Finally, when $\Gamma \frac{C}{q} < \bar{U}$, information-insensitive debt does not allow for borrowing the full pledgeable value $\bar{U}$, while information-sensitive debt is costly in terms of information production, having a credit capacity of $\bar{U} - C$. Now, the ranking between them is unclear. The social surplus generated by information-insensitive debt is $[(1-q)A - 1]\Gamma \frac{C}{q}$ which is greater than the social surplus generated by information-sensitive debt (or equity), $(1-q)A(\bar{U} - C) - \bar{U}$, if and only if

$$\frac{C}{\bar{U}} > \left[\frac{(1-q)A}{(1-q)A - 1} + \frac{1}{p\left(\frac{G-\bar{U}}{\bar{U}}\right) q} \right]^{-1}. \tag{4}$$

In words, information-insensitive debt dominates if the cost of acquiring information is sufficiently high relative to the expected pledgeable value, which is more likely to be satisfied when the project productivity, A, is high, when the liquidation value is not information-worthy, i.e., $p\left(\frac{G-\bar{U}}{\bar{U}}\right)$ is low, and importantly, when the probability of default (and the lender receiving the asset), q, is low.

The next proposition summarizes these results.

Proposition 1 (Pecking order) *The optimal financing structure has the following properties:*

1) *Information-sensitive debt and information-sensitive equity are equivalent.*

2) *Information-insensitive debt dominates information-insensitive equity when* $q < 1$.

3) *Information-insensitive debt dominates information-sensitive contracts under condition (4).*

Our model not only features a pecking order of financing structures but also explains why firms, and in particular financial firms, pledge assets as collateral and issue secured debt to cover liquidity needs (for example, through a repurchase agreement) rather than simply selling the assets. Selling assets is equivalent to selling equity shares in the asset, since in both cases the buyer or investor owns the asset outright. The incentive to examine the asset is stronger than that of a lender, who only receives the asset (collateral) when the borrower defaults. Secured debt reduces the incentive to produce costly information and thereby enlarges funding capacity.

Our analysis below shows that "lending no questions asked" should not be read as a slogan for indiscriminately suppressing information, but rather as a concrete principle for designing financial arrangements. It means structuring claims so that the benefit from privately examining assets remains below the point at which claimants find it worth the cost. In the simplest case, this favors debt over equity (or an outright asset sale). In more elaborate settings that we study below, it favors architectures that reduce collateral exposure link by link on intermediation chains by inserting intermediaries at funding bottlenecks, placing intermediaries with higher information costs upstream, and placing safer intermediaries downstream. The architecture matters because it determines how large a credit capacity and lending in the dark sustains.

Discussion: government bonds as collateral. Our setting interprets the collateral U as a firm's pledgeable value—capital equipment, real estate, or securities (e.g., mortgage-backed securities and other structured products)—whose value is uncertain at $t = 0$ and whose examination is costly. Government bonds are often used as collateral in rehypothecation chains in practice. We argue that the paper's mechanism extends to government bonds under a natural reinterpretation of the lender's information-acquisition decision.

A lender who receives a bond as collateral in the borrower's default may hold it rather than sell it for cash because cash offers inferior returns.[11] Given that the lender does not sell the bond, the bond's value to her depends jointly on the bond's payoff—driven by interest-rate and liquidity risk tied to macroeconomic conditions—and on the payoff of her existing assets, which load on the same macro shocks. For example, a lender whose own balance sheet is exposed to duration risk may care not about the bond's public, standalone market price, but about the bond's value in joint states in which her own funding position is weak. Under this reinterpretation, we have $U = G$ denote the favorable joint state and $U = B$ the unfavorable one.[12] The information cost C is the

[11]In payment systems, cash relaxes cash-in-advance constraints, so cash trades at a money premium, offering a lower equilibrium return than bonds. Cash also commands a regulatory premium, as it is treated more favorably under regulatory frameworks. A lender who does not face binding cash-in-advance or regulatory constraints therefore earns a lower return on cash than on bonds and strictly prefers to hold the bond rather than sell it for cash.

[12]Under this reinterpretation, the information cost reflects the lender's discovery of the collateral's private value, so U may differ between lender and borrower—warranting separate notation for the borrower's valuation. This distinction, however, does not affect our analysis: given sufficiently high A, the borrower's only objective is to maximize borrowing capacity, and her collateral value enters only through her net profits.

cost of learning this joint state—equivalently, the cost of learning the conditional distribution of her own asset's value given realized bond returns. Knowing this joint distribution is costly.

By paying the information cost, the lender refines her information set so that she can then back out of the financing contract in unfavorable joint states and commit to it in favorable states. This is the same temptation to deviate towards information acquisition that gives rise to the IC constraint limiting the capacity of information-insensitive debt. Here, the mechanism is applied to a joint state rather than a scalar collateral value.[13]

3 Credit Intermediation

We extend the environment by adding another firm that is ex-ante identical to the one considered in the previous section. To simplify the exposition, we assume that the two firms' liquidity shocks are uncorrelated. A firm's capacity to serve as an intermediary is unrelated to whether it needs liquidity itself. Its own liquidity needs can be met by pledging its project in a separate transaction.[14] What matters for a firm's intermediation capacity is its failure probability conditional on the end borrower's failure. We assume that the *unconditional probability* of failure is q for both firms, but given the failure of one firm, the *conditional probability* that the other firm fails is $\phi \in (0, 1)$.

When a firm is hit by a liquidity shock, the funds it can raise directly for lenders using information-insensitive debt is given by $\min\{\bar{U}, \Gamma C/q\}$ as shown in Lemma 2. We assume the condition (4) holds and focus on information-insensitive debt. The firm may alternatively seek intermediated financing — i.e., borrowing by pledging U as collateral to the other firm, which then becomes an intermediary and repledges the collateral to the end lender.[15]

Next, we show that intermediation enlarges the financing capacity of information-insensitive debt. But now, there are two financing links. One is *upstream*, which is the credit the end borrower obtains from the intermediary. The other is *downstream*, which is the credit the intermediary

[13]Because the lender's own assets differ in the exposure to bond risk factors, C is heterogeneous, which aligns with the heterogeneous-intermediary setup of Section 4. Our main mechanism then applies as it does under our baseline interpretation: intermediation dilutes each lender's probabilistic exposure to receiving the bond, discourages the costly production of information on the joint state of the bond and her own asset, and thereby enlarges credit capacity.

[14]The liquidity shock just determines which firm needs to borrow. Both firms can be hit by liquidity shocks, and they can still intermediate for each other, giving rise to two chains. If one firm is hit by a liquidity shock (becoming a borrower) and another is not, we have one chain in which the intermediary is the firm without borrowing needs. Once the borrowing needs are pinned down, the liquidity shock no longer shapes the chain's mechanics.

[15]The borrower need not observe the entire downstream chain of repledging to evaluate the option of intermediated financing. She compares the loan amount offered through direct financing with that offered by an intermediary. She chooses whichever option offers more credit. Her transaction with the lender is bilateral.

obtains from the end lender. In the downstream, the intermediary's debt capacity is given by

$$L_d = \min \left\{ \bar{U},\, \Gamma \frac{C}{q\phi} \right\}, \tag{5}$$

where the subscript "d" denotes "downstream". It is greater than the capacity of direct financing, $\min\{\bar{U}, \Gamma C/q\}$ in Lemma 2, because q is replaced by $q\phi$. Conditional on the end borrower's failure (with probability q), the intermediary fails with probability ϕ, in which case it lacks sufficient cash flows to repay the end lender and hands over the end borrower's collateral. Since the end lender gets the collateral when both firms fail, with the probability $q\phi$ (rather than q), it is less inclined to examine the end borrower's collateral when lending to the intermediary.

In words, intermediation introduces additional cash flows that buffer the end lender's risk of ending up with the end borrower's collateral. This is the first important role of intermediation: *by involving cash flows that are not perfectly correlated with the end borrower, intermediation creates ownership dilution—reducing the probability of collateral flowing downstream.* This effect relies on $\phi < 1$—the intermediary may succeed even if the end borrower fails.

In the upstream, the end borrower's information-insensitive debt capacity is

$$L_u = \min \left\{ \bar{U},\, \Gamma \frac{C}{q(1-\phi)} \right\}, \tag{6}$$

where the subscript "u" denotes "upstream". It is greater than the capacity of direct financing, $\min\{\bar{U}, \Gamma C/q\}$ in Lemma 2, because q is replaced by $q(1-\phi)$. The end borrower's failure does not imply that the intermediary will keep the collateral. The intermediary keeps the collateral with probability q and keeps it only if it survives, with conditional probability $1-\phi$. When the intermediary also fails, the collateral will be passed to the end lender. Since the intermediate lender's ultimate probability of *getting and keeping the collateral* is when the end borrower fails but the own project succeeds, $q(1-\phi)$, it is less inclined to examine the end borrower's collateral than the end lender in the case of direct financing without intermediation.

This is the second important role of intermediation: *because the intermediate lender's cash flows are risky and default can happen, intermediation enhances ownership dilution by reducing the probability that the collateral remains upstream*. This effect relies on $\phi > 0$—the intermediary project does not necessarily succeed if the end borrower's project fails.

Figure 2 represents these flows and the credit capacity with and without intermediation.

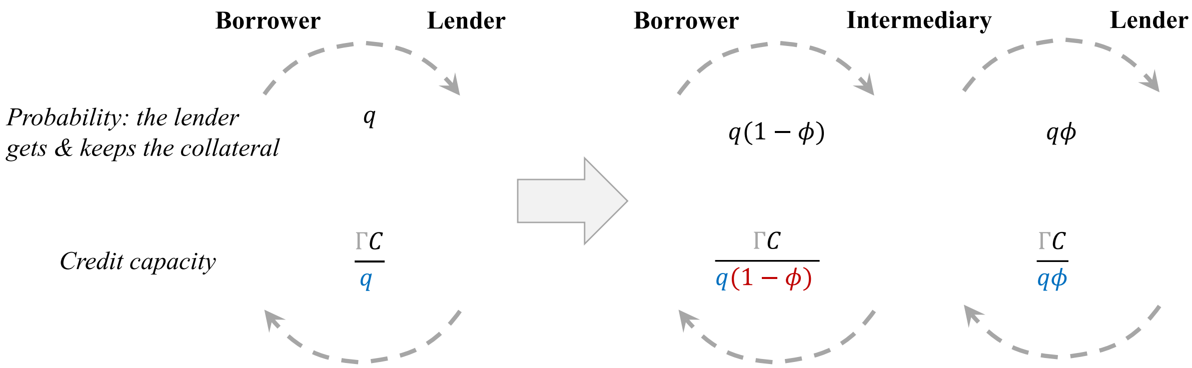

Figure 2: **From Direct Financing to Intermediated Financing.**

Intermediation ultimately increases the borrower's credit capacity, which is given by

$$L = \min\left\{\bar{U},\, L_d,\, L_u\right\} = \min\left\{\bar{U},\, \Gamma\frac{C}{q\phi},\, \Gamma\frac{C}{q(1-\phi)}\right\} \equiv \min\left\{\bar{U},\, \Gamma\frac{C}{\widehat{q}}\right\}, \tag{7}$$

where we define a composite probability

$$\widehat{q} = \max\{q\phi,\, q(1-\phi)\}. \tag{8}$$

When $\phi \in (0,1)$, $\widehat{q} < q$ and L is greater than credit with direct financing, $\min\{\bar{U}, \Gamma C/q\}$ in Lemma 2. From $\widehat{q} = \max\{q\phi,\, q(1-\phi)\}$, it is clear that intermediated financing boosts funding capacity by diluting *both* the end lender's and intermediary's incentive to produce information.

The composite probability, $\widehat{q} = \max\{q\phi,\, q(1-\phi)\}$, reflects a funding *"bottleneck"*. When $\phi > 1/2$, the bottleneck is located upstream, as the intermediary is more likely to end up with the collateral than the end lender. When $\phi < 1/2$, the bottleneck is located downstream, as the end lender is the most likely to end up with the collateral. Hence, the correlation structure that balances these forces and maximizes credit capacity is $\phi = 1/2$, i.e., when the two firms are uncorrelated—conditional on one firm's failure, the failure or success probability of the other is 50%. Under $\phi = 1/2$, *there is no bottleneck*, the dilution of asset ownership is maximized, and the end borrower's credit capacity is maximized. The next proposition summarizes this discussion.

Proposition 2 (Intermediated financing capacity) *The borrower's credit capacity with intermediation is given by (7), greater than that without intermediation, given by Lemma 2. Furthermore, credit capacity is maximized when the cash flows of the borrower's project and the intermediary's project are uncorrelated, i.e., $\phi = 1/2$, and there is no bottleneck.*

Several observations are in order. While we have shown that intermediation enlarges the capacity of information-insensitive debt, we have not established that information-insensitive debt is the optimal financing instrument under intermediation. Under condition (4), information-insensitive debt is optimal for direct financing, as shown in the previous section. Note that when A is sufficiently large (part of Assumption 1 in Section 2.1), the right side of (4) is increasing in q.[16] Since replacing q with $q\phi$ or $q(1-\phi)$ (both strictly less than q) reduces the right-hand side, the condition (4) continues to hold, and hence information-insensitive debt remains optimal under intermediation. In the next section, where we extend the intermediation chain to multiple intermediaries, we show more generally that the debt capacity of any link is $L = \min\{\bar{U}, \Gamma C/m\}$, where m is the probability that the lender on that link owns the collateral—necessarily smaller than q, as otherwise extending the chain would be redundant. Condition (4) therefore suffices to ensure that information-insensitive debt is optimal for intermediation chains of any length.

Intermediation requires the repledging of collateral, which can be *implicit* if collateral is transferred only upon default at $t = 2$ (usual bank credit), or *explicit* if the collateral is an asset that can be separated from the project's operation (e.g., a financial asset) and it is transferred along the credit chain at $t = 1$ and repurchased at $t = 2$ upon debt repayment (usual repo agreement with rehypothecation).[17] Under implicit repledging, collateral ownership transfers only upon the borrower's default, so the borrower faces no risk of losing her collateral. Under explicit repledging, however, the borrower is exposed to this risk: if her project succeeds but the intermediary becomes insolvent, the repledged collateral may be locked in the intermediary's bankruptcy process and not returned to the borrower (Infante, 2019; Infante and Vardoulakis, 2021). In Appendix B, we introduce such counterparty risk and analyze its impact on financing capacity and optimal chain length. We show that the chain remains beneficial and discuss the institutional safeguards that mitigate this risk and bring the setting close to our baseline model.[18] This extension yields additional empirical predictions: in environments with weaker collateral protection, credit chains should be shorter and rely on intermediaries whose solvency is more positively correlated with that of the borrower.

Our model distinguishes *collateral opaqueness* from *intermediation complexity* in their roles in increasing credit capacity. Asset opaqueness is about how costly it is to examine the asset, i.e.,

[16]Specifically, the right-hand side of (4), as a function of q, has an inverse-U shape, peaking at $q = (A-1)/(A+\sqrt{Ap(G-\bar{U})/\bar{U}})$, which is increasing in A.

[17]Intermediation in general can be viewed as rehypothecating claims (e.g., Rampini and Viswanathan (2019)).

[18]If the borrower repays, the intermediary has enough resources to repay the specific loan extended by the end lender, retrieve that collateral, and return it to the end borrower. This assumes, of course, that the loan is ring-fenced from the intermediary's other liabilities and the proceeds from the end borrower's repayment are segregated.

the information cost C, relative to how information-worthy the asset is, which is captured by the inverse of Γ in equation (2). In contrast, intermediation complexity concerns how the financing architecture dilutes asset ownership probabilistically along the chain. This is captured by $\widehat{q}$ in equation (8). While asset opaqueness has been studied in the literature (as reviewed by Gorton and Ordonez (2023)), our focus here is on intermediation complexity.

Finally, note that a firm acting as an intermediary does not preclude it from borrowing for its own liquidity needs through another intermediary. In fact, if both firms face liquidity needs, credit is maximized by using the other firm as an intermediary instead of each borrowing directly from the end lender. Along both credit chains, the end lender's exposure to either firm's default and associated collateral is buffered by the other firm's potential survival, weakening each end lender's incentives to provide information about the collateral of the other firm. Moreover, when acting as an intermediary, neither firm is guaranteed to own the other's collateral upon the other's default, because the intermediary firm itself may fail in that event. Here, intermediation does not add new resources. What increases credit capacity is the financial structure alone.

Expanding this logic to many firms, there may exist several intermediation chains that combine to form a complex network, with a seemingly spurious flow of funds between firms that make their own investments and also channel funds to one another, serving as both end borrowers and intermediate lenders on chains simultaneously. We explore these intricacies in the next section.

Discussion: chains vs. other arrangements. One alternative for diluting information-acquisition incentives is for the borrower to simultaneously obtain funds from multiple investors. In this case, each investor may have a small stake and a weak incentive to examine the asset. Such dilution differs from ours. First, it applies to equity as well as debt, whereas our theory highlights probabilistic asset ownership as a distinct advantage of debt and offers a rationale for debt chains. Second, if there is only one investor with funds (i.e., funds are concentrated), or if multilateral contracting is costly, having multiple investors becomes infeasible, but our intermediation chain can still be formed to probabilistically carve up asset ownership. Even an agent without available funds to lend, without superior technology, without superior information or expertise, but with only a risky balance sheet, can still join the chain as an intermediary and enlarge credit capacity.

The paper does not claim that chains dominate every conceivable contractual arrangement. The general logic of incentive compatibility in information production—the amount of financing that can be sustained is governed by the claimants' exposure that induces private examination—

applies to richer contracts that split the claimants' exposure to dilute their information-production incentive. The distinctive role of credit chains, which is our focus, is practical: they constitute the minimal bilateral architecture that implements this dilution. Our focus, in other words, is on the complex architecture of simple bilateral debt contracts, not on the design of complex contracts.

4 Intermediation Networks

Section 3 showed that introducing a single intermediary, by diluting the end lender's probabilistic ownership of the collateral, enlarges the end borrower's credit capacity. The logic extends naturally: adding a second intermediary can further dilute each lender's probabilistic ownership of the collateral, weakening her incentive to acquire information yet again. With multiple intermediaries available, two questions arise. *Normatively*, how is credit capacity maximized by the choice of a subset of intermediaries and ordering them on a chain? *Positively*, can the same chain emerge from decentralized decisions in the absence of any planner? We answer both questions in this section.

To capture an arbitrary chain length and full heterogeneity, we adopt the following notation. Index the end borrower as 0 and the end lender (the investor) as $I = n$; the integers $1, \ldots, n-1$ index intermediaries from upstream (closest to the end borrower) to downstream (closest to the end lender). Lender i's information cost is C_i. The probability that everyone upstream of i defaults—lender i's probability of *getting* the collateral—is γ^{i-1}, with the initial condition $\gamma^0 = q$. Conditional on everyone upstream defaulting, lender i's probability of survival—and thus of *keeping* the collateral—is $\theta_i^{i-1} = 1 - \phi_i^{i-1}$. The end lender does not default, so $\theta_n^{n-1} = 1$. We impose no parametric structure on the joint default distribution: the γ's and θ's may depend in any way on the identities of agents occupying upstream positions.

The lender of link i *gets* the collateral with probability γ^{i-1} and, conditional on getting it, *keeps* it with probability θ_i^{i-1}. Combining these two forces with the IC constraints of Section 3, the credit capacity of link i is $\Gamma C_i/[\gamma^{i-1}\theta_i^{i-1}]$, and the chain's credit capacity is determined by the bottleneck:

$$L = \min\left\{\bar{U},\, \Gamma \min_{i\in\{1,\ldots,n\}} \frac{C_i}{\gamma^{i-1}\theta_i^{i-1}}\right\}. \tag{9}$$

Sections 2 and 3 are special cases of (9), with no intermediary and one intermediary, respectively.

As a normative benchmark, one can ask which subset and ordering of intermediaries maximizes chain capacity. The positive question is whether a decentralized process of invitations and

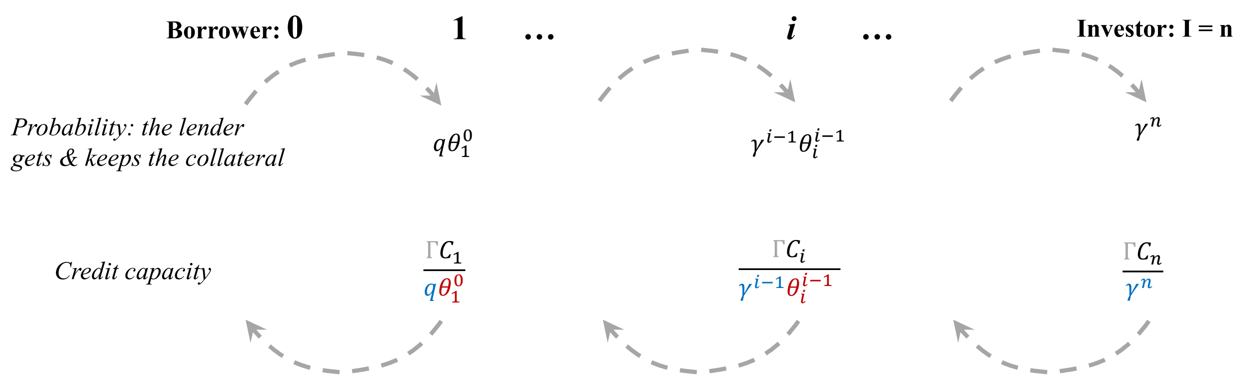

Figure 3: **A Chain with n Intermediaries.**

bilateral bargaining implements that same arrangement. Section 4 addresses both questions in sequence. First, given a number of participants, we define the capacity-maximizing chain, show that it exists, and that it can arise as the unique equilibrium outcome of decentralized formation. Second, we endogenize the number of participants. Finally, we characterize various properties of the socially optimal chain that arises as an equilibrium outcome.

4.1 Optimal Chain Formation

Let $i(\cdot)$ denote the ordering function: $i(k) = j$ means intermediary k occupies the j-th position on the chain. The optimal chain is defined as the following solution to the planner's problem:

$$\max_{\mathcal{S}, i(\cdot)} L\big(\mathcal{S}, \{i(k)\}_{k\in\mathcal{S}}\big), \tag{10}$$

where $\mathcal{S}$ is the set of chain participants and, given $\mathcal{S}$ and the participants' positions $\{i(k)\}_{k\in\mathcal{S}}$, L is defined in (9). The optimization is defined over a finite set of candidates and their ordering, so a planner with full information about candidate participants, their information costs, and their joint default distribution can, in principle, solve it directly. We first record that an optimal chain exists.

Lemma 3 (Optimal chain existence) *The candidate set is finite, so the set of all possible chain participants and their ordering on a credit chain is finite. And, given that the credit capacity of any* (9) *is bounded above by $\bar{U}$, a chain that attains the maximum credit capacity therefore exists.*

The shared setup: a sequential game. To study decentralized chain formation, we consider the following sequential game. An *initiator*—one of the two endpoints of the chain (i.e., the end borrower and the investor)—extends an invitation to a candidate intermediary with the outside option being to transact directly with the opposite endpoint. The inviter and invitee then bargain over the surplus generated by the resulting link via Nash bargaining. The invitee's outside option is normalized to zero. After the link is formed, the intermediary may reach out to the next invitee. In the paper, endogenous intermediation refers to a chain arising from decentralized invitations, participation decisions, and bilateral bargaining.

At any step of chain formation, the inviter's disagreement payoff is determined by the chain capacity she would secure by terminating the chain—an action defined on the chain assembled so far and independent of the identity of the candidate currently at the bargaining table. The following protocol specifies the off-equilibrium disagreement payoff: at any step k, the inviter selects a single candidate to negotiate with, and disagreement terminates chain extension—if the current inviter is the opposite endpoint, the chain ends here; otherwise, the current inviter reaches out to the opposite endpoint.[19] On the subgame-perfect equilibrium (SPE) path, invited candidates enlarge chain capacity, resulting in a Pareto improvement, so bargaining never fails if such candidates exist. The invitee's disagreement payoff is zero. If the invitee accepts, she becomes the next inviter and extends an invitation to another candidate, with the same bargaining structure. The game continues until an inviter chooses to terminate the chain by connecting to the opposite endpoint.

The identity of the initiator is exogenously specified—either the end borrower 0 or the investor I—and the initiator is the residual claimant of the chain surplus that accrues through all the bilateral Nash-bargaining outcomes along the chain; this is why she has an incentive to initiate the chain.[20] A SPE exists and can be computed by backward induction. The two theorems below consider the two natural choices of initiator—the end borrower and the investor (end lender).

Borrower-initiated formation. Our baseline specification of Section 2 assigns the residual claim to the end borrower; she therefore has the strict incentive to initiate the chain, and we consider first a borrower-initiated chain. The next theorem characterizes the resulting SPE.

[19] An alternative protocol of "switch-to-next-best" specification would let the inviter's disagreement payoff be tied to the chain capacity under her next-best candidate. Then disagreement payoff depends on which candidate is at the table (because "next-best" is defined relative to the pool with the current candidate excluded), and for sufficiently low Nash weight the inviter may have an incentive to bargain with a suboptimal candidate in order to sharpen her threat point against the best one—a hold-up wrinkle that termination-only sidesteps. The SPE chain composition is, in any case, invariant under any monotone surplus-sharing rule, as we note below; only the surplus distribution differs across rules.

[20] Endogenizing the identity of the initiator requires a model of search and matching, beyond the scope of this paper.

Theorem 1 (Borrower-initiated chain formation) *Suppose the end borrower* 0 *is the initiator of the sequential game and the residual claimant of the chain surplus, as in Section 2. A unique SPE exists. On the equilibrium path, every invitee accepts, and the resulting chain attains the maximum credit capacity of Lemma 3 and coincides with the planner's solution. Each member of the chain earns a strictly positive expected payoff.*

The proof, in Appendix A.6, proceeds by backward induction. At each step, the current inviter's Nash-bargaining payoff decomposes into two parts: her disagreement payoff from the outside option—the chain capacity she could secure by terminating the chain now and connecting directly to the end lender—and her share of the marginal surplus that the current candidate adds on top of the outside option. The first part does not depend on which candidate is currently at the bargaining table, because termination is defined solely on the existing chain, irrespective of the candidate being negotiated with. The second part is monotone increasing in the chain capacity from the game continuation generated by the invitee. Hence, the inviter's best response is simply to invite the candidate that maximizes the chain's capacity—the same objective a planner would solve at that step in a recursive formulation. Propagating this logic backward along the SPE path, the sequential game implements the chain that coincides with the planner's optimum.

Two features of the equilibrium deserve emphasis. First, the rents are shared all the way up the chain. Each intermediary collects Nash-bargaining surplus both as the invitee on her own link and, if she is not the last to invite, as the inviter on the next link. The joint surplus telescopes upward to the residual claimant—the end borrower, under this theorem—but each intermediary still captures strictly positive rents from her dual role. This is the economic source of the unique-SPE claim: no candidate wants to deviate off the equilibrium path, because accepting an optimal invitation and extending optimally strictly dominates either rejecting or inviting suboptimally.[21] Second, no separate commitment device is required. The bargaining at each link is bilateral, and the resulting contract is binding; the inviter's disagreement payoff is fully determined by her outside option—terminating the chain at that step and connecting directly to the opposite endpoint—which is defined regardless of any party's commitment power. The SPE therefore does not rely on the contracts at different links being negotiated simultaneously or conditionally on each other.

[21] Other monotone surplus-sharing rules—such as the Shapley value, or a rule that gives all bargaining power to the inviter—generate the same chain composition because, in each case, the inviter's payoff remains monotone in continuation chain capacity. How surplus is distributed, of course, differs across rules.

Lender-initiated formation. Borrower-initiated chain formation is a natural benchmark, but many real lending relationships originate on the lender side. A money market fund actively searches for short-term counterparties to deploy its cash; an insurance company sources investment-grade structured-credit positions; a bank extends loans to business development companies. The intermediaries that receive the end lenders' funds in turn seek borrowers. The next theorem shows that a symmetric construction, in which the chain is built from the lender side, implements the same optimal chain—provided the surplus assignment is adjusted accordingly, so that the lender has a strict incentive to initiate. Specifically, we consider the following: the end lender I is the initiator of the sequential game and the residual claimant of the chain surplus (a modification of the baseline specification of Section 2, in which the end borrower is the residual claimant). The sequential game now runs upstream: I invites a candidate who will borrow from her, that candidate invites a further upstream candidate, and so on, until the chain reaches the end borrower 0.

Theorem 2 (Lender-initiated chain formation) *Suppose the end lender (investor) I is the initiator of the sequential game and the residual claimant of the chain surplus. A unique SPE exists. On the equilibrium path, every invitee accepts, and the resulting chain attains the maximum credit capacity of Lemma 3 and coincides with the planner's solution. Each intermediary on the chain earns a strictly positive expected payoff.*

The proof, in Appendix A.7, is symmetric to that of Theorem 1. The only change in the model setup is the identity of the residual claimant: reassigning the residual surplus from the end borrower to the end lender gives the latter an incentive to initiate the chain. The same recursive incentive-alignment argument and backward induction delivers the capacity-maximizing chain.

Discussion: information environments. Theorems 1 and 2 are structurally parallel but differ in where the informational burden falls. Under borrower-initiated formation (Theorem 1), the end borrower must solve the backward-induction problem from her side in order to extend the first invitation. That means she must know the full candidate pool and the joint default structure: her choice of first intermediary depends on anticipating which intermediary the first invitee will, in turn, invite, which the second will then invite, and so on, all the way down to the end lender. Each subsequent inviter in the chain has the same foresight, but the end borrower bears the greatest informational burden because her choice determines the entire downstream chain.

Under lender-initiated formation (Theorem 2), the informational burden sits on the opposite endpoint. The end lender and each subsequent inviter must know the candidate pool to plan the chain backward from the lender's side. The end borrower, by contrast, is a passive participant: by the time the chain reaches her, the most upstream intermediary approaches with a single loan offer, and the borrower's only decision is to accept or decline. She does not need to know who sits downstream of that intermediary, because the downstream chain enters her decision problem only through the loan terms she is quoted.

This asymmetry has empirical content. When the end borrower has limited visibility into the financing landscape—as is typical for a household borrower in a mortgage chain or a small business tapping a private credit fund—lender-initiated formation is the natural institutional arrangement, and Theorem 2 is the relevant positive benchmark. When the end borrower has dedicated treasury expertise and broad visibility into the financing market—as is plausible for large corporate borrowers structuring their own financing chains along with financial advisors—borrower-initiated formation is also plausible, and Theorem 1 applies. The two theorems together show that the ownership-dilution mechanism is robust to the question of who bears the informational burden: whichever side has the capacity to plan the chain can implement the capacity-maximizing arrangement, and the associated SPE delivers strictly positive rents to every participating intermediary.

4.2 Chain Extension

Section 4.1 assumes the set of potential participants is fixed and asks how candidates are invited onto a chain and in what order. We now ask what happens to an existing chain when new candidates become available: should a new participant be admitted? It is evident that, should a new intermediary be added, she must be inserted into the bottleneck link: while inserting an intermediary into the non-bottleneck link may enlarge the credit capacity of that link, the credit capacity of the whole chain remains determined by the link with the smallest capacity for channeling funds.

Suppose the current chain has $n-1$ intermediaries and a bottleneck at link $(i, i+1)$ where a new intermediary can be inserted. What criteria should the new intermediary meet for the extended chain to achieve a greater credit capacity? For this, it is important to highlight an asymmetry in the location of an intermediary within a chain. A new intermediary does not change anything for those "upstream" but changes both the probabilities of getting and of keeping the asset downstream. While it is clear that adding a new intermediary reduces the probability that downstream intermediaries receive the asset (γ), its effect on the probability that downstream agents retain the asset (θ) is not obvious, since this depends on the correlation structure of participants' projects in the chain.

We introduce a natural assumption: θ—an intermediary's probability of project success conditional on the end borrower's failure and all upstream intermediaries' failure—is lower when there are more upstream intermediaries. This is intuitive: surviving is more difficult when more agents fail in the economy, likely because the economic environment is worse. Given that adding a new intermediary reduces both the *probability of getting the asset* and the *probability of keeping the asset* for those downstream, our focus can be solely on comparing the lending capacity of the current bottleneck lender with that of the new intermediary. The next proposition summarizes this condition for chain extension. The proof is in Appendix A.8.

Proposition 3 (Chain extension condition) *If the current intermediation chain with $n-1$ intermediaries has not maximized financing capacity to the full pledgeable value $\bar{U}$, the addition of the n-th intermediary to the bottleneck $(i, i+1)$, $i \leq n-1$, enlarges credit capacity if and only if*

$$\frac{C_n}{\theta_n^i} > \frac{C_{i+1}}{\theta_{i+1}^i}. \tag{11}$$

When comparing the current bottleneck lender (the i+1-th intermediary) and the new intermediary, we should consider three elements: the information cost, the probability of getting the asset, γ, and the probability of keeping the asset, θ. In the inequality (11), γ drops out because the new intermediary takes the same position as the old, and the probability of getting the asset just depends on the position in the chain, not the identity of the intermediary. In contrast, information costs and the probability of keeping the asset do depend on the intermediary's identity. The condition has an intuitive interpretation: inserting the new intermediary enlarges credit capacity (widens the bottleneck) if its information cost is higher or its probability of keeping the asset conditional on getting the asset is lower. Both forces cause the new intermediary to have a weaker incentive to examine the asset than the current bottleneck lender.

4.3 Characterizing the Optimal Chain

Sections 4.1 and 4.2 characterize the SPE chain and the marginal condition for adding a new intermediary at the bottleneck. Both results apply to any candidate pool – finite or continuous, sparse or dense. We now add a *richness* assumption on the pool: candidate types vary continuously, so the planner can fine-tune the type at each link rather than being limited to a discrete menu. Under richness, the optimal chain admits a tight characterization: every link operates at the same per-link capacity, the highest sustainable bottleneck level. This equalization condition then yields clean sorting criteria when we further assume conditional independence of intermediaries' defaults.

The candidate pool's primitives are the information cost C (a candidate-level characteristic) and the joint default distribution across candidates (a chain-dependent characteristic that is the probability of getting and keeping the collateral, i.e., γ and θ, at each position). The pool satisfies *richness* if C and, for any chain, γ and θ are continuously distributed over a bounded but open set. [22] Under the additional assumption of conditional independence introduced in Corollary 1 below, an intermediary k's conditional survival probability $\theta(k)$ is itself a primitive, and richness reduces to the condition that (C, θ) is continuously distributed on a bounded open set in $\mathbb{R}^2$.

With continuous control over candidate types, the planner can re-balance incentives across links: if some link's lender faces less temptation than the bottleneck lender, the "incentive slack" should be reallocated. The optimum equalizes the IC tightness and credit capacity across all links. Note that with a discrete and coarse candidate pool, exact equalization may not be feasible.

Proposition 4 (Incentive Equalization) *Under richness, the optimal chain equalizes link capacities at the bottleneck level: for every intermediary k on the optimal chain at position $i^*(k)$,*

$$L^* \;=\; \Gamma \frac{C(k)}{\gamma^{i^*(k)-1}\theta_{i^*(k)}^{i^*(k)-1}}, \tag{12}$$

where L^* is the optimal credit capacity.

The proof, in Appendix A.9, is a reallocation argument. Consider a slack at any link (i.e., the link's credit capacity being greater than that of the bottleneck), in which case the capacity equalization condition fails. We show that this leads to a contradiction. If the bottleneck is downstream of this slack link, marginally raising the slack-link lender's conditional survival probability θ tightens her own link's capacity (toward L^*) but lowers γ at downstream positions, strictly raising downstream link capacities, including the bottleneck. If the bottleneck is upstream, directly lowering the bottleneck candidate's θ raises her own link, increasing γ for the downstream, which is acceptable given the slack. In both cases, the perturbation strictly increases the chain's capacity, contradicting the chain's optimality. At the optimum, no such profitable reallocation exists, which then implies that at the optimum, all link capacities are equal and coincide at L^*. The intuition is akin to output maximization under the Leontief production function.

[22] The open-set part is essential: every candidate in the pool is interior, so marginal perturbations of any primitive in any direction yield other candidates in the pool, and the planner is never blocked at a boundary.

The equalization condition delivers a clean characterization of the optimal chain under increasingly sharp assumptions on the joint default distribution. First impose *conditional independence*: an intermediary k's probability of survival conditional on the end borrower's default and any subset of upstream defaults equals a marginal probability $\theta(k)$ that depends only on k:

$$\theta(k) \equiv \mathbf{Prob}(k \text{ survives} \mid 0 \text{ defaults}) = \mathbf{Prob}(k \text{ survives} \mid 0 \text{ and any other intermediary defaults}).$$

Note that, here, k indexes intermediaries, not their positions on the chain. Under this condition, we have $\gamma^{j-1} = q \prod_{i(k)<j}(1-\theta(k))$, and the condition (12) yields the following single-index sort.

Corollary 1 (Sort by the C/θ ratio) *Under conditional independence, the optimal chain orders intermediaries by the ratio $C(k)/\theta(k)$ in decreasing order from upstream to downstream.*

The intuition is straightforward. At any link, the credit capacity is given by the lender's $C/(\gamma\theta)$, where γ, the probability of getting the collateral (upstream joint failure), is decreasing in the lender's position, as the number of middlemen between her and the end borrower increases. Therefore, to balance the decreasing γ in the denominator, C/θ must be decreasing as well. Given $\theta(k)$, an intermediary k with a high $C(k)$ is reluctant to acquire information at any position, so the optimal chain can afford to place her upstream and give her a high probability of getting the collateral—k's upstream joint default probability is high because there are few intermediaries between k and the end borrower. A high-$\theta(k)$ candidate is unlikely to pass collateral downstream, making her a natural "safety valve" close to the end lender, where upstream defaults have already diluted the probability of her getting collateral, and her role is to act as a buffer for her downstream.

When candidates are homogeneous along one of the two dimensions, this ratio reduces to a sort along the other dimension.

Corollary 2 (Optimal sequencing under homogeneous θ) *Under $\theta(k) = \theta \in (0,1)$, an intermediary k with a lower information cost is placed downstream, i.e., $i^*(k)$ is higher.*

Corollary 3 (Optimal sequencing under homogeneous C) *Under $C(k) = C$, an intermediary k with a higher conditional survival probability $\theta(k)$ is placed downstream, i.e., $i^*(k)$ is higher.*

These corollaries formalize the empirical pattern noted in the introduction: downstream intermediaries are safer in two senses. By their position, they face a reduced probability of upstream default—more upstream intermediaries reduce their probability of getting the collateral—and, by the optimal sequencing, they also have higher conditional survival probability due to their asset and business characteristics. For example, along the chains in private credit markets, banks and money-market funds sit downstream from BDCs and private credit funds, as the corollaries predict.

Discussion: distributing asset ownership vs. risk. Forming intermediation chains introduces additional randomness—each intermediary's survival is not perfectly correlated with others'. In our setting, that randomness is at the heart of intermediation benefits: it carves up the probabilistic asset ownership (the first moment of receiving the collateral) and thereby dilutes the incentive within the chain to examine that collateral. While the additional randomness may pile up higher moments, agents in our model are risk-neutral, so higher moments do not enter the decision problem. Typically in the literature, financial networks are formed to distribute risk; in our model, they are formed to distribute ownership. The risk-neutrality assumption isolates the latter mechanism.

5 Specialized Funding Intermediaries

So far, our analysis has focused on firms' liquidity needs at $t = 1$, conditional on firms investing K in the projects. Next, we consider firms' asset choices at $t = 0$. Firms either invest K in the project or in a bond. The bond is simply a storage technology—for each unit of investment at $t = 0$, the firm receives one, so its return is $r^b = 1$. Under this normalization, absent frictions, no firm has incentives to invest in the bond at $t = 0$. As we will show, however, in our setting the bond is useful to facilitate intermediation, hence generating a rent from improving credit capacity, a sort of *convenience yield* that justifies their ubiquitous presence within intermediation networks.

To make this point clear, we go back to the setting with two firms. Besides a lower return, investing in bonds differs from investing in projects in two ways. First, investing in bonds makes it more difficult to examine projects, with a cost $\bar{C} > C$. This means to capture that, by specializing in an investment, the firm relinquishes expertise in examining other firms' assets involved in projects. Second, bonds have lower returns, but they are certain: there is symmetric information about the bond's value. This makes bonds free from informational concerns when used as collateral. In what follows, we will study under what conditions having a *bond-holding intermediary* is beneficial, the value of its services, and the conditions under which a firm chooses bonds over projects.

A firm that invests in bonds can borrow up to $K > \bar{U}$ (the face value of bonds with no information concerns) from the end lender. If a firm in need of liquidity borrows *directly* from the end lender, it can obtain $L^d = \frac{\gamma C}{q}$ issuing information-insensitive debt (see Section 2). If borrowing instead from a *bond-holding* intermediary, it can obtain $L^b = \frac{\gamma \bar{C}}{q} > L^d$.

The larger credit capacity translates into an expected social surplus at $t = 0$:

$$S^b = \lambda[(1-q)A - 1](L^b - L^d) = \lambda[(1-q)A - 1]\Gamma\frac{(\bar{C} - C)}{q}, \tag{13}$$

where λ is the probability of liquidity shock and financing needs, and, as discussed in Section 2, $(1-q)A - 1$ is the expected return from investing at $t = 1$ in the liquidity event.

Let h denote the fraction of surplus seized by the bond-holding intermediary firm.[23] We can compute the expected return of the bond-holding intermediary at $t = 0$:

$$r^b = \frac{K + hS^b}{K} = 1 + \lambda h[(1-q)A - 1]\left(\frac{\Gamma}{K}\right)\frac{(\bar{C} - C)}{q}, \tag{14}$$

The bond return has two components: storage technology and intermediation profit. Let V^d denote the value of a project under direct financing (i.e., the project-investing firm's outside option) in the liquidity event rather than financing intermediated by the bond-holding intermediary. The project-investing firm's return is

$$r^p = \frac{V^d + (1-h)S^b}{K}.$$

The return also has two components. A baseline return V^d/K under direct financing in the liquidity event and $1 - h$ fraction of the additional surplus from intermediation facilitated by bonds.[24]

In equilibrium, firms must be indifferent between investing in bonds and projects, $r^b = r^p$ or

$$h = \frac{1}{2} + \left(\frac{V^d - K}{2S^b}\right). \tag{15}$$

In addition, for the equilibrium to exist, the bond return must be sufficiently high such that the bond-investing firm does not deviate to investing in a project, i.e., $r^b \geq V^i/K$, or equivalently,

$$h \geq \frac{V^i - K}{S^b}, \tag{16}$$

[23] In Section 3, the surplus created by intermediation was captured by the end borrower (i.e., $h = 0$). Here, we consider a more general case where the surplus can be split between the borrower and the intermediary.

[24] The expected project value under direct financing is given by $V^d = \bar{U} + \lambda\left[(1-q)A - 1\right]L^d + (1-\lambda)(1-q)AK$, which, as shown in Figure 1, consists of a liquidation value, $\bar{U}$, across all event branches, investment in the liquidity event at $t = 1$, and in the absence of liquidity shock, the expected value of a success project of scale K.

where V^i is the project value under intermediated financing.[25] Note that when the bond-holding firm deviates, both firms invest in projects and intermediate for each other as in Section 3. Finally, we verify that the project-investing firm does not deviate to holding the bond. This is obvious: deviation means both firms holding the bond and no need for intermediation at $t = 1$, which implies $r^b = 1$ below the project return, which is greater than 1 as we have assumed that the project productivity, A, is sufficiently high.

Proposition 5 (Intermediation facilitated by bonds) *The equilibrium with one firm investing in the bond and the other in a project exists with h given by (15) if the condition (16) holds.*

Which type of intermediation dominates? The social value of both firms investing in projects is $2V^i$. The social value of a single firm investing in the project is $K + V^d + S^b$: even though the economy misses the value of the project, the bond-holding intermediary facilitates intermediation and thus increases reinvestment in the event of liquidity needs. If $K + V^d + S^b > 2V^i$, the equilibrium with less ex-ante investment in positive NPV projects but more efficient intermediation dominates.

Our model demonstrates that two types of financial architecture can arise under costly information production. In Sections 3 and 4, our focus is on intermediation that dilutes the probabilistic ownership of *collateral of uncertain quality*, which firms take as given. In this section, we show that allowing asset choices at $t = 0$ enlarges credit capacity: bond-holding intermediaries, instead of diluting probabilistic asset ownership, serve as information barriers by replacing *collateral of uncertain quality* with *collateral of certain quality*, such as government bonds. This result suggests that even more complex networks exist, combining different types of collateral assets, with uncertain high returns and certain low returns. This extension shows how foregoing ex-ante profitable investments can be rationalized by enhancing the economy's credit capacity.

6 Conclusion

When not enough is at stake, an investor is willing to lend "in the dark" rather than examine the pledged asset, at an information cost ultimately borne by the borrower and reducing her financing capacity. Equity of the pledged asset exposes the investor to the asset under all circumstances, whereas debt represents a probabilistic ownership: the lender owns the collateral only upon default.

[25]The project value is given by $V^i = \bar{U} + \lambda\left[(1-q)A - 1\right] L^i + (1-\lambda)(1-q)AK$.

Debt, therefore, dominates equity in sustaining information-insensitive financing capacity. This pecking order is not our main contribution; the optimality of debt has been established in prior studies, albeit in different settings. What is new is that characterizing debt as a probabilistic claim on the pledged asset, and the implications of this property for the lender's incentive to produce information, allow us to develop a new theory of financial intermediation and credit chains.

Intermediation chains arise endogenously to dilute the probabilistic asset ownership that debt contracts represent. For intermediaries further downstream and the end lender, the likelihood of ultimately acquiring the collateral declines with distance from the original borrower, a form of ownership dilution along the chain. Intermediaries positioned upstream in the optimal chain are correlated in financial health with the end borrower and with intermediaries preceding them, so their effective ownership of the collateral is diluted by the possibility that, upon receiving it, they may default and pass the collateral downstream rather than retain it. Therefore, seemingly spurious chains can be rationalized as a means of expanding credit capacity. This result constitutes a cautionary tale for regulators who aim to enhance transparency by reducing the layers of fund flows, as they must also weigh the unintended consequences, including reduced financing capacity.

Our mechanism is most relevant in environments where claimants can profit from privately examining collateral and where such examination is costly enough to matter for financing capacity. In those settings, the model predicts that intermediation chains should be organized around information incentives: bottleneck links should attract additional intermediation, safer or more diversified balance sheets should appear downstream, and chain length should be limited by intermediary homogeneity. Where these informational forces are weak, other mechanisms—counterparty risk, collateral scarcity, safe-asset demand, or maturity transformation—are likely to become the dominant factors behind chain formation. The paper's contribution is to isolate the conditions under which the architecture of credit intermediation is itself a source of credit creation. The optimal architecture depends on intermediaries' information costs and the chain-specific joint failure distribution. In addition, we show that intermediation needs may induce certain firms to specialize in investing in assets, such as government bonds, that in principle have lower returns than productive projects but are valuable for an intermediary to "grease" the credit flow along the chain.

Our model can be applied to understand various forms of debt contracts and intermediation, for example, repurchase agreements (repo) and rehypothecation, which are essentially mechanisms for diluting the end and intermediate lenders' incentives to produce costly information. Repo requires the borrower's asset to be transferred to the lender on the spot and returned upon repayment of the debt, while other debt contracts channel funds to the borrower on the spot and transfer the asset only in bankruptcy. The timing of collateral transfer is inconsequential in our model.

References

Acharya, V. V., M. Gopal, M. Jager, and S. Steffen (2025). Shadow always touches the feet: Implications of bank credit lines to non-bank financial intermediaries. Working Paper 33590, National Bureau of Economic Research.

Atkeson, A. G., A. L. Eisfeldt, and P. Weill (2015). Entry and exit in otc derivatives markets. *Econometrica 83*(6), 2231–2292.

Babus, A. (2016). The formation of financial networks. *The RAND Journal of Economics 47*(2), 239–272.

Babus, A. and P. Kondor (2018). Trading and information diffusion in over-the-counter markets. *Econometrica 86*(5), 1727–1769.

Brumm, J., M. Grill, F. Kubler, and K. Schmedders (2023). Re-use of collateral: Leverage, volatility, and welfare. *Review of Economic Dynamics 47*, 19–46.

Brusco, S. and F. Castiglionesi (2007). Liquidity coinsurance, moral hazard, and financial contagion. *Journal of Finance 62*(5), 2275–2302.

Caballero, R. J. and A. Simsek (2013). Fire sales in a model of complexity. *The Journal of Finance 68*(6), 2549–2587.

Chang, B. and S. Zhang (2015). Endogenous market making and network formation. Working paper, London School of Economics and Political Science.

Chernenko, Sergey, R. I. and D. Scharfstein (2025). Bank capital and the growth of private credit. Working paper.

Craig, B. and Y. Ma (2022). Intermediation in the interbank lending market. *Journal of Financial Economics 145*(2, Part A), 179–207.

Dang, T. V., G. Gorton, B. Holmström, and G. Ordoñez (2017). Banks as secret keepers. *American Economic Review 107*(4), 1005–29.

Diamond, D. (1984). Financial intermediation and delegated monitoring. *Review of Economic Studies 51*, 393–414.

Diamond, D. W. (1985). Optimal release of information by firms. *The Journal of Finance 40*(4), 1071–1094.

Diamond, D. W. and R. E. Verrecchia (1991). Disclosure, liquidity, and the cost of capital. *The Journal of Finance 46*(4), 1325–1359.

Donaldson, J. R. and E. Micheler (2018). Resaleable debt and systemic risk. *Journal of Financial Economics 127*(3), 485–504.

Donaldson, J. R., G. Piacentino, and X. Yu (2022). Systemic risk in financial networks revisited: The role of maturity. Working paper, University of Colorado Boulder and University of Southern California.

European Commission (2022). Simplified securitization framework in the eu.

Farboodi, M. (2023). Intermediation and voluntary exposure to counterparty risk. *Journal of Political Economy 131*(12), 3267–3309.

Federal Reserve Bank of Dallas (2012). The too-big-to-fail problem. *Dallas Fed Economic Letter*.

Federal Reserve Bank of New York (2020). Bank complexity and financial stability. *FRBNY Staff Reports* (880).

G20 Financial Stability Board (2018). Post-crisis financial reforms: Achievements and next steps.

Glode, V. and C. Opp (2016). Asymmetric information and intermediation chains. *American Economic Review 106*(9), 2699–2721.

Glode, V., C. C. Opp, and X. Zhang (2019). On the efficiency of long intermediation chains. *Journal of Financial Intermediation 38*, 11–18.

Gorton, G. and G. Ordonez (2023). *Macroeconomics and Financial Crises: Bound Together by Information Dynamics*. Princeton university Press.

Gottardi, P., V. Maurin, and C. Monnet (2019). A theory of repurchase agreements, collateral re-use, and repo intermediation. *Review of Economic Dynamics 33*, 30–56. Fragmented Financial Markets.

Haque, S., Y. S. Jang, and J. J. Wang (2025). Indirect credit supply: How bank lending to private credit shapes monetary policy transmission. Working paper, Board of Governors of the Federal Reserve System.

He, Z. and J. Li (2022). Intermediation via credit chains. Working Paper 29632, National Bureau of Economic Research.

Hendershott, T., D. Li, D. Livdan, and N. Schürhoff (2020). Relationship trading in over-the-counter markets. *The Journal of Finance 75*(2), 683–734.

Hirshleifer, J. (1971). The private and social value of information and the reward to inventive activity. *The American Economic Review 61*(4), 561–574.

Hugonnier, J., B. Lester, and P.-O. Weill (2019). Frictional intermediation in over-the-counter markets. *The Review of Economic Studies 87*(3), 1432–1469.

Infante, S. (2019). Liquidity windfalls: The consequences of repo rehypothecation. *Journal of Financial Economics 133*(1), 42–63.

Infante, S., C. Press, and J. Strauss (2018). The ins and outs of collateral re-use. Feds notes, Board of Governors of the Federal Reserve System (U.S.).

Infante, S. and Z. Saravay (2024). Treasury re-use and the demand for safe assets. Working paper, Board of Governors of the Federal Reserve System.

Infante, S. and A. P. Vardoulakis (2021, 12). Collateral runs. *The Review of Financial Studies 34*(6), 2949–2992.

Jiang, E. X. (2023). Financing competitors: Shadow banks' funding and mortgage market competition. *The Review of Financial Studies 36*(10), 3861–3905.

Maggio, M. D. and A. Tahbaz-Salehi (2014). Collateral shortages and intermediation networks. Working paper, Imperial College and Northwestern University.

Monnet, C. and B. N. Narajabad (2012). Why rent when you can buy? A theory of repurchase agreements. 2012 Meeting Papers 647, Society for Economic Dynamics.

Myers, S. and N. Majluf (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics 13*, 187–221.

Parlatore, C. (2019). Collateralizing liquidity. *Journal of Financial Economics 131*, 299–322.

Prudential Regulation Authority (2022). Uk bank ring-fencing: Review and recommendations.

Rampini, A. A. and S. Viswanathan (2019). Financial Intermediary Capital. *The Review of Eco-*

nomic Studies 86(1), 413–455.

Rock, K. (1986). Why new issues are underpriced. *Journal of Financial Economics 15*(1), 187–212.

Sannino, F. (2026). Committing to trade: A theory of intermediation. *Journal of Financial Economics 178*, 104249.

Stiglitz, J. (1999). Reforming the global economic architecture: Lessons from recent crises. *Journal of Finance 54*, 1508–1521.

Viswanathan, S. and J. J. D. Wang (2004). Inter-dealer trading in financial markets. *The Journal of Business 77*(4), 987–1040.

A Proofs

A.1 Proof of Lemma 1

A.1.1 Information-sensitive equity

The firm offers an investor a contract specifying the equity price of k_U^s for z_U^s of the equity shares, where $U \in \{B, G\}$ and the superscript denotes the type of financing contract ("s" for information-sensitive). After the investor acquires information at the cost of C, the value of U is revealed to her and to the firm, and the U-contingent contract is executed. For the investor to participate, the break-even condition must hold:

$$pz_G^s(G - k_G^s - C) + (1-p)z_B^s(B - k_B^s - C) = 0. \tag{A.1}$$

It is assumed that the investor breaks even in expectation, and all the surplus goes to the firm.

Given Assumption 1 (i.e., A is sufficiently high), the optimal fraction of share issuance is $z_G^s = z_B^s = 1$. The values of E_B^s and E_G^s hence represent the financing the firm can obtain to continue operations if the liquidation value is low or high, respectively. Given the linearity of the constraint, these are indeterminate, so we normalize $E_B^s = B$, which simply implies that the cost of information acquisition is ex-ante compensated if the liquidation value is G. The investor's break-even condition (A.1) implies $k_G^s = G - C/p$. In words, given a sufficiently high A, the firm is willing to sell the liquidation value that is worth G at a discounted price $k_G^s = G - C/p$ so that the investor breaks even in expectation, taking into account the cost of information production. Note that the cost of information production reduces the firm's financing capacity: when $U = G$, the funds raised are k_G^s, which is below G (the pledgeable liquidation value of the firm's stock).

We assume the firm keeps the full surplus, which is equal to

$$(1-q)\left[p(Ak_G^s - G) + (1-p)(Ak_B^s - B)\right] + q(0 - \bar{U}), \tag{A.2}$$

where $Ak_G^s - G$, for instance, captures the profits from obtaining funds by selling the asset with high liquidation value and continuing operations at a scale E_G^s, with a nonpledgeable return A in case the project succeeds, with probability $1-q$. This scenario happens with probability p and the corresponding situation with low liquidation value with probability $1-p$. If the project fails, the firm loses the liquidation value to the investor. Substituting $E_B^s = B$ and $k_G^s = G - C/p$ into the social surplus, we obtain the following result.

Social surplus. The optimal information-sensitive equity contract generates social surplus $[(1-q)(A-1)-q]\bar{U}-(1-q)AC$. The first part of the surplus, $[(1-q)(A-1)-q]\bar{U}$, shows that it is increasing in the expected pledgeable value, $\bar{U}$. Consistent with our assumption of a sufficiently high A, we maintain $(1-q)(A-1)-q>0$ throughout the paper. The second term, $-(1-q)AC$, shows that the cost of information production reduces the surplus by reducing financing capacity and wasting resources on producing information. The cost is higher when the project is more productive, i.e., A is higher, and when it is more likely to succeed, i.e., $1-q$ is higher.

A.1.2 Information-insensitive equity

The firm offers to sell equity at a price E^i for a z^i fraction of the liquidation (pledgeable) value, where the superscript, "i", represents information-insensitive. Notice that in this case, there is no subscript as the contract is by construction not conditional on the true liquidation value. For this contract to be feasible, the investor must not have an incentive to deviate and privately learn about U at a cost of C. The equity price, E^i, is set so that the investor breaks even based on the expected liquidation value:

$$E^i = z^i\bar{U} = z^i[pG+(1-p)B]. \tag{A.3}$$

The investor does not privately produce information if its expected return(the left side below) is lower than the expected return of following the contract without information acquisition (zero profit on the right side below):

$$(1-p)(0-C)+p(z^iG-E^i-C)\leq 0. \tag{A.4}$$

On the left side, the first term represents the case of $U=B$, where the investor will not buy equity as the payout is below the price, $z^iB<E^i=z^i\bar{U}$, and the second term represents the gain of knowing $U=G$ privately but buying equity at the lower uncertainty price. Rearranging this incentive-compatibility (IC) condition and substituting out E^i using (A.3), we obtain

$$E^i = z^i\bar{U} \leq \frac{C}{p\left(\frac{G-\bar{U}}{\bar{U}}\right)}, \tag{A.5}$$

where the right side is a limit on the amount of information-insensitive equity financing the firm can raise. Intuitively, if the information cost is low or the liquidation (pledgeable) value is information worthy (i.e., the percentage deviation of G from $\bar{U}$ is high), the investor is tempted to produce information, so the information-insensitive financing capacity is low. We consolidate such properties of the pledgeable value into one parameter, Γ, given by (2), so that the IC constraint on financing capacity can be written as

$$E^i = z^i \bar{U} \leq \Gamma C, \tag{A.6}$$

where Γ summarizes the attributes of the pledgeable value that induce information and C is the cost of such information. The lower Γ and C, the higher the incentives to learn about the asset, and the stronger the constraint in raising funds with an information-insensitive equity contract.

The surplus from this contract, which is the information-insensitive counterpart of (A.2), is

$$(1-q)(AE^i - z^i \bar{U}) + q(0 - z^i \bar{U}), \tag{A.7}$$

When $E^i = \bar{U}$ (i.e., the IC constraint (A.6) is not binding at $z^i = 1$), the social surplus is greater than that given by (A.2) under information-sensitive equity financing. However, the maximum financing capacity $\bar{U}$ may not be attainable under the IC constraint (A.6).

Social surplus. The optimal information-insensitive equity contract is subject to the IC constraint (A.6) that limits z^i, the fraction of equity sold to the investor. Given z^i, it generates social surplus $[(1-q)(A-1)-q]z^i \bar{U}$ in the liquidity event. Note that as long as the investor's break-even condition holds, the information-sensitive contract is feasible. The information-insensitive contract requires the additional IC condition (A.6). ■

A.2 Proof of Lemma 2

A.2.1 Information-sensitive debt

The timing is the same as in the equity contract case. The firm offers a contract that is U-contingent. After receiving the contract, the lender produces information on U, and U is also revealed to the firm. Then the contract is executed under $U = B$ or $U = G$. Information-sensitive debt is feasible as long as the lender's break-even condition holds.

The debt contract is summarized by three variables. The lending amount is denoted by L^s_U, the repayment to the lender is denoted by R_U^s, and the fraction of collateral or liquidation value the lender seizes in default is denoted by x_U^s, where $U \in \{B, G\}$. In the following, we characterize the optimal contract step-by-step.

Given the assumption that A is sufficiently high, and that the firm wants to borrow as much as possible, we set $x_B^s = x_G^s = 1$. This is, the firm wants to pledge as much collateral as possible.

Whether the project fails or not is the firm's private information, so the firm may default even if it succeeds.[1] Therefore, a debt contract must induce the firm to tell the truth and hand the collateral over only when it cannot repay. The truth-telling condition requires $R_G^s = G$ and $R_B^s = B$. Note that for the equity contract, we do not discuss this issue because whether the project fails or succeeds, the lender receives the same payoff given by the liquidation (pledgeable) value.

The lender's break-even (participation) condition equates the expected loan minus the information costs with the expected repayment and default proceedings.

$$pL_G^s + (1-p)L_B^s + C = (1-q)[pR_G^s + (1-p)R_B^s] + q[px_G^s G + (1-p)x_B^s B]. \tag{A.8}$$

Replacing full collateral pledged ($x_B^s = x_G^s = 1$) and truth-telling conditions ($R_G^s = G$ and $R_B^s = B$), we can rewrite the participation constraint as :

$$pL_G^s + (1-p)L_B^s + C = pG + (1-p)B = \bar{U}. \tag{A.9}$$

The payments in the two states, $U = B$ and G, are indeterminate, so we fix $L_B^s = B$, and obtain

$$pL_G^s + C = pG, \text{ or, equivalently, } L_G^s = G - \frac{C}{p}. \tag{A.10}$$

Therefore, we have fully characterized the U-contingent debt contract (L_G^i, R_G^i, x_G^i) and (L_B^i, R_B^i, x_B^i). For $U = G$, we have $L_G^s = G - \frac{C}{p}$, $R_G^s = G$, and $x_s^G = 1$, and for $U = B$, $L_B^s = B$, $R_B^s = B$, and $x_s^F = 1$. Even though collateral is fully pledged, the expected credit capacity is below the expected liquidation value by the cost of information C.

[1] Note that verifying project outcome is not an issue for equity financing because, whether the project succeeds or fails, the equity investor's share is predetermined.

Social surplus. The social surplus created by the information-sensitive debt is given by

$$p(1-q)\left[AL_G^s - (1-q)R_G^s - qx_G^s G\right] + (1-p)(1-q)\left[AL_B^s - (1-q)R_B^s - qx_B^s B\right]. \quad \text{(A.11)}$$

The optimal information-sensitive debt contract generates social surplus $[(1-q)(A-1)-q]\bar{U} - (1-q)AC$ in the liquidity event.

A.2.2 Information-insensitive debt

Next, we consider the scenario in which the lender lends based on the expected liquidation value, and it is incentive-compatible for the lender not to produce costly information. Without information about the liquidation value, U, the debt contract is no longer contingent on U. It specifies the amount of lending, L^i (the superscript "i" is for information-insensitive), the fraction of liquidation value seized by the lender when the project fails, x^i, and the nominal repayment, R^i. The lender's break-even (participation) condition is,

$$L^i = (1-q)R^i + qx^i[pG + (1-p)B], \quad \text{(A.12)}$$

and the firm's truth-telling condition,

$$R^i = x^i[pG + (1-p)B]. \quad \text{(A.13)}$$

From the borrower's truth-telling condition, we obtain

$$x^i = \frac{R^i}{pG + (1-p)B} = \frac{R^i}{\bar{U}} \leq 1, \quad \text{(A.14)}$$

where the last inequality captures the fact that the firm cannot pledge more than the entire collateral. Substituting this solution of x^i into the lender's participation (break-even) condition, we obtain

$$L^i = R^i. \quad \text{(A.15)}$$

After the firm proposes the debt contract, the lender decides whether to produce information and whether to accept the offer. Therefore, the contract design is subject to the following incentive compatibility constraint:

$$0 \geq (1-p)(0-C) + p[(1-q)R^i + qx^iG - L^i - C]. \tag{A.16}$$

The left side represents the case without information production, as the lender breaks even and earns zero profit. On the right side, if $U = B$, the lender will not lend to the firm as the expected payoff is smaller than the specified lending amount, as shown below:

$$qR^i + (1-q)x^iB = L^i \left[q + (1-q)\frac{B}{\bar{U}} \right] < L^i,$$

where we apply (A.15) to substitute out R^i and x^i in the first step If $U = G$, the lender will accept the offer, generating positive profits. The positive profits under $U = G$ are directly implied by the break-even in expectation and the loss from lending under $U = B$. Thereby, we have confirmed that under $U = B$, the lender declines the offer, and under $U = G$, the lender accepts the contract.

The incentive compatibility (A.16) constraint can be simplified to the following inequality that is at the heart of our model and carries several key messages:

$$L^i \leq \frac{C}{qp\left(\frac{G-\bar{U}}{\bar{U}}\right)} = \Gamma\frac{C}{q}, \tag{A.17}$$

where we use Γ, defined in (2), to summarize the attributes of the project's collateral. As with equity, the left side (financing capacity) is high when it is costly for the lender to acquire information (high C) or the asset does not induce information (high Γ). Importantly, however, in a debt contract, credit capacity is also high when the probability of the project failing is low (low q). This is not due to the lender's risk aversion and credit risk being priced in equilibrium. In our model, the lender is risk-neutral. The link between default probability and credit capacity emerges from the lender's information choice. When q is low, it is unlikely that the firm will default and the lender will end up in possession of the asset, discouraging its examination.

Since A is sufficiently high, the constraint (A.17) binds, and we have fully characterized the information-insensitive debt contract, (L^i, R^i, x^i):

$$L^i = \min\left\{\bar{U}, \Gamma\frac{C}{q}\right\}, \qquad R^i = L^i, \quad \texttt{and} \quad x^i = R^i/\bar{U}. \tag{A.18}$$

Social surplus. The social surplus, i.e., the firm's profit, is

$$(1-q)(AL^i - R^i) + q(0 - x^i\bar{U}) = [(1-q)(A-1) - q]L^i. \tag{A.19}$$

The optimal information-insensitive debt contract is subject to the IC constraint (A.17) that limits L^i, the amount of lending. Given L^i, it generates surplus $[(1-q)(A-1)-q]L^i$ in the liquidity event. ■

A.3 Proof of Proposition 1

From Lemmas 1 and 2, and equations (3) and (4). ■

A.4 Proof of Proposition 2

From comparing equations (7) and (8). ■

A.5 Proof of Lemma 3

The candidate set is finite, so the set of (subset of chain participants and their ordering) pairs is finite. The credit capacity of chain σ, denoted by $L(\sigma)$, in (9) is well-defined for each pair and bounded above by $\bar{U}$. A chain that attains the maximum credit capacity, therefore exists. ■

A.6 Proof of Theorem 1

The game is finite: the candidate set $\mathcal{N}$ is finite, no agent appears twice on a chain, so the chain length is bounded by $|\mathcal{N}|+1$; at each node the action set {terminate, invite a candidate or terminate}, i.e., terminate $\cup$ $(\mathcal{N} \setminus \sigma_{<k})$ where $\sigma_{<k}$ is a chain with $k-1$ intermediaries (positions $0, 1, \ldots, k-1$ occupied), is finite; and given continuation values, the bilateral Nash bargaining at each link admits a unique solution as we show below. By the Zermelo–Kuhn result, a subgame perfect equilibrium (SPE) of pure strategies exists and can be computed by backward induction.

By Lemma 2, a fully assembled chain σ ending at I generates social surplus

$$S(\sigma) \;\equiv\; \big[(1-q)A - 1\big]\, L(\sigma) \;=\; m\, L(\sigma),$$

where $L(\sigma)$ is the chain capacity from (9) and $m \equiv (1-q)A - 1 > 0$ is the per-unit-capacity surplus multiplier (positive whenever the project is NPV-positive). Let $V(\sigma_{<k})$ denote the SPE continuation surplus at the assembled prefix $\sigma_{<k}$ (the borrower-side portion of the chain assembled so far)—the surplus the chain will ultimately generate when all subsequent moves follow the SPE. The recursive formulation is

$$V(\sigma_{<k}) \;=\; \max\Big\{\, S\big(\sigma_{<k} \cup \{I\}\big),\; \max_{c \in \mathcal{N} \setminus \sigma_{<k}} V\big(\sigma_{<k} \cup \{c\}\big) \,\Big\}, \tag{A.20}$$

where the first term inside the outer max is the surplus from terminating the chain at step k by connecting directly to the end lender (investor) I, and the second term is the continuation surplus from extending the chain by inviting some candidate c. The recursion ends when no candidate remains, or further extension is unprofitable.

Consider the inviter at step k (the agent at position $k-1$) deciding which action to take. We adopt the natural specification that her disagreement payoff in the bilateral Nash problem with candidate c is the surplus under termination at step k, $V^{\text{out}}(\sigma_{<k}) \equiv S\big(\sigma_{<k} \cup \{I\}\big)$, which does not depend on c. The invitee's disagreement payoff is zero. The marginal Nash surplus generated by including candidate c is then

$$\Delta_k(c \mid \sigma_{<k}) \;\equiv\; V\big(\sigma_{<k} \cup \{c\}\big) \;-\; V^{\text{out}}(\sigma_{<k}),$$

and the inviter's total surplus payoff under Nash bargaining with weight $\alpha \in (0, 1]$ is

$$u_k^{\text{inv}}(c) \;=\; V^{\text{out}}(\sigma_{<k}) \;+\; \alpha\, \Delta_k(c \mid \sigma_{<k}) \;=\; (1-\alpha)\, V^{\text{out}}(\sigma_{<k}) \;+\; \alpha\, V\big(\sigma_{<k} \cup \{c\}\big).$$

Since $V^{\text{out}}(\sigma_{<k})$ does not depend on c, the inviter's surplus payoff $u_k^{\text{inv}}(c)$ is strictly increasing in the continuation surplus $V(\sigma_{<k} \cup \{c\})$: she always prefers the candidate delivering the highest continuation surplus. Augmenting the action set with a pseudo-candidate "terminate" whose continuation value is defined as $V(\sigma_{<k} \cup \{\text{terminate}\}) \equiv S\big(\sigma_{<k} \cup \{I\}\big)$, the inviter's optimal action is simply the one that maximizes the continuation surplus $V(\sigma_{<k} \cup \{c\})$ over $c \in \{\text{terminate}\} \cup (\mathcal{N} \setminus \sigma_{<k})$. This is exactly the surplus-maximizing recursion (A.20).[2]

Recursion (A.20) has two interpretations that turn out to coincide. Read as an *SPE recur-*

[2]The argument extends to other monotone surplus-sharing rules—Shapley value, inviter-takes-all, etc.—because in each case the inviter's payoff remains monotone in $\tilde{V}(\{c\} \cup \tilde{\sigma})$.

sion, $V(\sigma_{<k})$ is the SPE continuation surplus starting from the assembled prefix $\sigma_{<k}$. Read as a *planner's recursion*, $V(\sigma_{<k})$ is the value of the planner's surplus-maximization sub-problem, $\max_\tau S\big((\sigma_{<k} \cup \tau) \cup \{I\}\big) = m \max_\tau L\big((\sigma_{<k} \cup \tau) \cup \{I\}\big)$, over all continuations τ that extend the prefix $\sigma_{<k}$. Because $m > 0$ is a constant common to all chains, surplus-maximization and capacity-maximization yield the same optimal chain. Bellman's principle of optimality—any tail of an optimal chain must itself be optimal given the assembled prefix—guarantees that the planner's value function satisfies the same recursion (A.20), with $V(\{0\}) = \max_\sigma S(\sigma) = m\,L^* \equiv S^*$; existence of L^* (and hence of S^*) is guaranteed by Lemma 3.

The argument above established that, at every node $\sigma_{<k}$ along the game tree, the SPE inviter selects an action in $\arg\max_c V(\sigma_{<k} \cup \{c\})$—precisely the action the planner would take at the same node when solving (A.20). Since the SPE and the planner play identical actions at every node, a straightforward backward induction shows that the two realize identical values $V(\sigma)$ at every prefix σ: at any prefix where the argmax is to terminate, both players terminate and $V(\sigma) = S(\sigma \cup \{I\})$ is pinned down directly; at any prefix where the argmax is to extend, both invite the same $c^* \in \arg\max_c V(\sigma \cup \{c\})$ and $V(\sigma) = V(\sigma \cup \{c^*\})$, which coincides under the two readings by induction. Applying this to the borrower's initial prefix $\{0\}$ yields $V(\{0\}) = S^* = m\,L^*$, so the SPE chain attains the planner's optimum (in both surplus and capacity terms).

Finally, we show that every chain member earns a strictly positive total surplus. For the invitee at step k on the SPE path, the marginal Nash surplus $\Delta_k(c_k^{SPE} \mid \sigma_{<k})$ is strictly positive because the SPE chooses to extend (i.e., the SPE choice c_k^{SPE} satisfies $V(\sigma_{<k} \cup \{c_k^{SPE}\}) > V^{\text{out}}(\sigma_{<k})$); otherwise the chain would terminate at step k, and this invitee would not be on the chain. Hence, the invitee receives the strictly positive Nash surplus share $(1-\alpha)\,\Delta_k > 0$. By the same argument, every chain member who invites the next member also earns the strictly positive Nash surplus share on the next link as the inviter. The last intermediary on the chain (the one that connects to I) earns only the invitee-side surplus, which is strictly positive by the argument above. Hence, every chain member's total surplus is strictly positive, and her participation constraint is strictly satisfied. ■

A.7 Proof of Theorem 2

The proof mirrors that of Theorem 1, with the chain built upstream from the end lender I rather than downstream from the end borrower 0. Under the modified surplus assignment, the end lender is the residual claimant of chain surplus and therefore has a strict incentive to initiate the chain.

The game is finite by the same argument as in Theorem 1: the candidate set $\mathcal{N}$ is finite, no agent appears twice on a chain, the action set at each node is finite, and given continuation values, the bilateral Nash bargaining at each link admits a unique solution. By the Zermelo–Kuhn result, an SPE in pure strategies exists and can be computed by backward induction.

Let $\tilde{\sigma}$ denote a partial chain assembled from the lender side—a subset of $\{I\} \cup \mathcal{N}$ that includes I as its most-downstream node, plus zero or more upstream intermediaries already invited (at the initiation $\tilde{\sigma} = \{I\}$). Let $\tilde{V}(\tilde{\sigma})$ denote the SPE continuation surplus at this suffix (the lender-side portion of the chain assembled so far). As in the proof of Theorem 1, the chain surplus is $S(\sigma) \equiv m\, L(\sigma)$ with $m \equiv (1-q)A - 1 > 0$ from Lemma 2. The recursive formulation is

$$\tilde{V}(\tilde{\sigma}) \;=\; \max\Big\{\, S\big(\{0\} \cup \tilde{\sigma}\big),\; \max_{c \in \mathcal{N} \setminus \tilde{\sigma}} \tilde{V}\big(\{c\} \cup \tilde{\sigma}\big) \,\Big\}, \tag{A.21}$$

where the first term inside the outer max is the surplus from terminating the chain by having the most-upstream agent in $\tilde{\sigma}$ lend directly to the end borrower 0, and the second term is the continuation surplus from extending the chain upstream by inviting some candidate c (who becomes the new most-upstream node).

Consider the inviter at the current step (the most-upstream agent in $\tilde{\sigma}$, which is I at the start of the game). Her disagreement payoff in the bilateral Nash problem with candidate c is the surplus under termination, $\tilde{V}^{\text{out}}(\tilde{\sigma}) \equiv S\big(\{0\} \cup \tilde{\sigma}\big)$, which does not depend on c. The invitee's disagreement payoff is zero. The marginal Nash surplus from including candidate c is

$$\tilde{\Delta}(c \mid \tilde{\sigma}) \;\equiv\; \tilde{V}\big(\{c\} \cup \tilde{\sigma}\big) \;-\; \tilde{V}^{\text{out}}(\tilde{\sigma}),$$

and the inviter's total surplus payoff under Nash bargaining with weight $\alpha \in (0, 1]$ is

$$u^{\text{inv}}(c) \;=\; (1-\alpha)\, \tilde{V}^{\text{out}}(\tilde{\sigma}) \;+\; \alpha\, \tilde{V}\big(\{c\} \cup \tilde{\sigma}\big).$$

Since $\tilde{V}^{\text{out}}(\tilde{\sigma})$ does not depend on c, the inviter's surplus payoff is strictly increasing in the continuation surplus $\tilde{V}(\{c\} \cup \tilde{\sigma})$. Augmenting the action set with a pseudo-candidate "terminate" whose continuation value is defined as $\tilde{V}(\{\text{terminate}\} \cup \tilde{\sigma}) \equiv S\big(\{0\} \cup \tilde{\sigma}\big)$, the inviter's optimal action is the one that maximizes the continuation surplus $\tilde{V}(\{c\} \cup \tilde{\sigma})$ over $c \in \{\text{terminate}\} \cup (\mathcal{N} \setminus \tilde{\sigma})$. This is exactly the surplus-maximizing recursion (A.21).

Recursion (A.21) has two interpretations that coincide. Read as an *SPE recursion*, $\tilde{V}(\tilde{\sigma})$ is the continuation surplus under SPE play starting from the assembled suffix $\tilde{\sigma}$. Read as a *planner's recursion*, $\tilde{V}(\tilde{\sigma})$ is the value of the planner's surplus-maximization sub-problem, $\max_\tau S\big(\{0\} \cup \tau \cup \tilde{\sigma}\big) = m \max_\tau L\big(\{0\} \cup \tau \cup \tilde{\sigma}\big)$, over all upstream extensions τ of the suffix $\tilde{\sigma}$. Because $m > 0$ is a constant common to all chains, surplus-maximization and capacity-maximization yield the same optimal chain. Bellman's principle of optimality—any head of an optimal chain must itself be optimal given the assembled suffix—guarantees that the planner's value function satisfies the same recursion (A.21), with $\tilde{V}(\{I\}) = \max_\sigma S(\sigma) = m\, L^* \equiv S^*$; existence of L^* (and hence of S^*) is guaranteed by Lemma 3.

The argument above established that, at every node $\tilde{\sigma}$ along the game tree, the SPE inviter selects an action in $\arg\max_c \tilde{V}(\{c\} \cup \tilde{\sigma})$—precisely the action the planner would take at the same node when solving (A.21). Since the SPE and the planner play identical actions at every node, a straightforward backward induction shows that the two realize identical values $\tilde{V}(\tilde{\sigma})$ at every suffix $\tilde{\sigma}$: at any suffix where the argmax is to terminate, both players terminate and $\tilde{V}(\tilde{\sigma}) = S(\{0\} \cup \tilde{\sigma})$ is pinned down directly; at any suffix where the argmax is to extend, both invite the same $c^* \in \arg\max_c \tilde{V}(\{c\} \cup \tilde{\sigma})$ and $\tilde{V}(\tilde{\sigma}) = \tilde{V}(\{c^*\} \cup \tilde{\sigma})$, which coincides under the two readings by induction. Applying this to the lender's initial suffix $\{I\}$ yields $\tilde{V}(\{I\}) = S^* = m\, L^*$, so the SPE chain attains the planner's optimum (in both surplus and capacity terms).

Finally, we show that every chain member earns a strictly positive total surplus. For the invitee at any step on the SPE path, the marginal Nash surplus $\tilde{\Delta}(c^{SPE} \mid \tilde{\sigma})$ is strictly positive because the SPE chooses to extend (i.e., the SPE choice c^{SPE} satisfies $\tilde{V}(\{c^{SPE}\} \cup \tilde{\sigma}) > \tilde{V}^{\text{out}}(\tilde{\sigma})$); otherwise the chain would terminate at this step, and this invitee would not be on the chain. Hence, the invitee receives the strictly positive Nash surplus share $(1-\alpha)\, \tilde{\Delta} > 0$. By the same argument, every chain member who, in turn, extends the chain further upstream, additionally earns the strictly positive Nash surplus share on the next link as the inviter. The most upstream intermediary (the one who lends directly to 0) earns only the invitee-side surplus, which is strictly positive by the argument above. Hence, every chain member's total surplus is strictly positive, and her participation constraint is strictly satisfied. ■

A.8 Proof of Proposition 3

Inserting the n-th intermediary between i and i+1 enlarges financing capacity if and only if

$$\Gamma \min\left\{\frac{C_1}{\gamma^0\theta_1^0}, \frac{C_2}{\gamma^1\theta_2^1}, ..., \frac{C_i}{\gamma^{i-1}\theta_i^{i-1}}, \frac{C_n}{\gamma^i\theta_n^i}, \frac{C_{i+1}}{\gamma^i\phi_n^i\hat{\theta}_{i+1}^i}, ..., \frac{C_{n-1}}{\hat{\gamma}^{n-2}\hat{\theta}_{n-1}^{n-2}}, \frac{C_I}{\hat{\gamma}^{n-1}\hat{\phi}_n^{n-1}}\right\} >$$
$$\Gamma \min\left\{\frac{C_1}{\gamma^0\theta_1^0}, \frac{C_2}{\gamma^1\theta_2^1}, ..., \frac{C_i}{\gamma^{i-1}\theta_i^{i-1}}, \quad \frac{C_{i+1}}{\gamma^i\theta_{i+1}^i}, \ ..., \ \frac{C_{n-1}}{\gamma^{n-2}\theta_{n-1}^{n-2}}, \frac{C_I}{\gamma^{n-1}\phi_n^{n-1}}\right\}$$

where $\hat{\gamma}^i$ represents the joint probability of the borrower 0's default, the first i intermediaries' default, *and* the new intermediary's default, $\hat{\phi}_{i+1}^i$ is the probability of Intermediary i+1's default conditional on borrower 0's default, the first i intermediaries' default, *and* the newly inserted intermediary's default, and $\hat{\phi}_{i+1}^i$ is the probability of Intermediary i+1's survival conditional on borrower 0's default, the first i intermediaries' default, *and* the newly inserted intermediary's default. As a reminder of our notation, we have, by definition, $\gamma^{n-1}\phi_n^{n-1} = \gamma^n$ and $\hat{\gamma}^{n-1}\hat{\phi}_n^{n-1} = \hat{\gamma}^n$.

This condition is essentially an IC constraint on the network's expansion, as we need the total surplus of those already on the chain to improve. An important insight is that inserting a new intermediary changes these probabilities for those "downstream" in the chain, but not "upstream." Since the left and right sides share the initial i items, the condition is equivalent to

$$\Gamma \min\left\{\frac{C_n}{\gamma^i\theta_n^i}, \frac{C_{i+1}}{\gamma^i\phi_n^i\hat{\theta}_{i+1}^i}, ..., \frac{C_{n-1}}{\hat{\gamma}^{n-2}\hat{\theta}_{n-1}^{n-2}}, \frac{C_I}{\hat{\gamma}^{n-1}\hat{\phi}_n^{n-1}}\right\} > \Gamma \min\left\{\frac{C_{i+1}}{\gamma^i\theta_{i+1}^i}, ..., \frac{C_{n-1}}{\gamma^{n-2}\theta_{n-1}^{n-2}}, \frac{C_I}{\gamma^{n-1}\phi_n^{n-1}}\right\}. \tag{A.22}$$

Note that, for any $k \in \{1, ..., n-(i+1)\}$, we have, by the definitions of probabilities,

$$\frac{C_{n-k}}{\hat{\gamma}^{n-k-1}\hat{\theta}_{n-k}^{n-k-1}} > \frac{C_{n-k}}{\gamma^{n-k-1}\theta_{n-k}^{n-k-1}},$$

because the events measured by $\hat{\gamma}_{n-k}^{n-k-1}$ contains one more event (the new intermediary's default) than the events measured by γ_{n-k}^{n-k-1} so $\hat{\gamma}_{n-k}^{n-k-1} \le \gamma_{n-k}^{n-k-1}$. Next, we impose the following assumption:

$$\hat{\theta}_{n-k}^{n-k-1} \le \theta_{n-k}^{n-k-1}. \tag{A.23}$$

The probability for Intermediary n-k to survive is lower when another intermediary (the new one) defaults. By definition, $\hat{\theta}_{n-k}^{n-k-1}$ measures the probability of survival conditional on all the preceding

$n-k-1$ intermediaries' default *and* the new intermediary's default, while θ_{n-k}^{n-k-1} measures the probability of survival conditional on only the previous $n-k-1$ intermediaries' default.

An interesting property emerges: when a new intermediary is inserted into the chain, both γ and θ decrease. The former is lower because it is a joint probability, and from γ to $\hat{\gamma}$, another event is added (the new intermediary's default). The latter is lower as it is a conditional probability of survival, so according to our assumption, when more intermediaries default (i.e., adding the new intermediary's default), the economic environment is likely worse, so θ declines to $\hat{\theta}$.

Therefore, in the condition for the newly inserted intermediary to enlarge financing capacity, i.e., the inequality (A.22), from the third terms on the left side to the last term, they are all larger than the corresponding terms on the right side (i.e., the second to the last terms). Moreover, since the edge (i, i+1) is the bottleneck before we insert the n-th intermediary, we know that the right side can be simplified to just $\frac{C_{i+1}}{\gamma^i\theta_{i+1}^i}$. On the left side, we can ignore the third-to-last terms as they are larger than the second-to-last terms on the right side. So, the condition (A.22) can be simplified to

$$\min\left\{\frac{C_n}{\gamma^i\theta_n^i}, \frac{C_{i+1}}{\gamma^i\phi_n^i\hat{\theta}_{i+1}^i}\right\} > \frac{C_{i+1}}{\gamma^i\theta_{i+1}^i}, \tag{A.24}$$

where we also divide both sides by Γ. Note that $\frac{C_{i+1}}{\gamma^i\phi_n^i\hat{\theta}_{i+1}^i} \geq \frac{C_{i+1}}{\gamma^i\theta_{i+1}^i}$ because $\phi_n^i \in [0,1]$. Therefore, for the inequality to hold, we only need

$$\frac{C_n}{\theta_n^i} > \frac{C_{i+1}}{\theta_{i+1}^i}, \tag{A.25}$$

where, to simplify the expression, we divide both sides by γ^i. ■

A.9 Proof of Proposition 4

Setup. Let D_k denote the default event of the agent at position k, with D_0 the borrower's default. Let $L_j(\sigma) \equiv \Gamma C_j/[\gamma^{j-1}\theta_j^{j-1}]$ denote the lending capacity at link j in chain σ, where

$$\gamma^{j-1} \equiv \Pr(D_0 \cap D_1 \cap \cdots \cap D_{j-1}) \quad \text{and} \quad \theta_j^{j-1} \equiv \Pr(\neg D_j \mid D_0 \cap \cdots \cap D_{j-1})$$

are derived from the joint default distribution of chain participants—a primitive of the model that we do *not* assume to factorize into marginals. The chain rule gives, for any $m > j$,

$$\gamma^{m-1} = \gamma^{j-1} \prod_{k=j}^{m-1} (1 - \theta_k^{k-1}), \tag{A.26}$$

which holds unconditionally. The chain capacity is $L(\sigma) = \min\{\bar{U}, \min_j L_j(\sigma)\}$ from (9); we focus on the interior case $L^* < \bar{U}$, where the bottleneck binds at some intermediary's link. Then $L^* \equiv \max_\sigma \min_j L_j(\sigma)$ is the maximum feasible bottleneck lending capacity; the maximum exists by Lemma 3. In the paper's convention, position 0 is the end borrower, position n is the end lender, and positions increase toward the lender; we refer to higher-position links as *downstream* (closer to the end lender) and lower-position links as *upstream* (closer to the borrower). For clarity of exposition, we prove the proposition assuming the bottleneck at the optimum is unique; the multi-bottleneck case is discussed at the end.

The candidate pool satisfies richness if, for any chain σ, any link j on σ, and any small $\varepsilon > 0$, there exists a substitute candidate j' for the agent at position j such that:

(i) Information cost is unchanged: $C_{j'} = C_j$.

(ii) The conditional survival probability can be moved in either direction: $\theta_{j'}^{j-1} = \theta_j^{j-1} \pm \varepsilon$.

(iii) The downstream conditional survival probabilities are preserved: θ_k^{k-1} is unchanged for every $k > j$.

Property (iii) is the substantive content of richness under general dependence: for any link, the pool must contain candidates whose joint default behavior matches the original at the downstream-correlation dimension while differing in own-conditional-survival probability. Because the pool's primitives lie in a bounded *open* set, every candidate—in particular, the one currently at position j—is interior, so substitutes meeting (i)–(iii) exist for any small $\varepsilon > 0$ in *either* direction. Under conditional independence (the special case of Corollary 1), θ_k^{k-1} depends only on agent k's primitives—not on j—so (iii) is automatic; under general dependence, (iii) requires the pool to span a continuum of joint default behaviors along the appropriate direction.[3]

Take any optimal chain σ with $\min_j L_j(\sigma) = L^*$. We show that, under richness, no link can have strict slack. The argument is the planner's optimality condition: slack at any link means

[3] A weaker but sufficient version of (iii) replaces "unchanged" with "changes by $o(\varepsilon)$," so that the leading-order effect on γ^{m-1} for $m > j$ comes entirely from the direct change in $(1 - \theta_j^{j-1})$ in (A.26).

the planner is providing more incentive than the bottleneck demands, and that incentive can be redistributed to relax the binding constraint.

Suppose for contradiction that $L_j(\sigma) > L^*$ at some position j, with the unique bottleneck L^* attained at $b \neq j$. We split into two cases based on the chain positions of j and b.

Case 1: $b > j$ **(bottleneck downstream of the slack link).** Substitute the candidate at position j with one satisfying $\theta_{j'}^{j-1} = \theta_j^{j-1} + \varepsilon$, $C_{j'} = C_j$, and θ_k^{k-1} unchanged for $k > j$.

Own link. L_j is proportional to $1/\theta_j^{j-1}$, so L_j decreases by $\approx L_j\,\varepsilon/\theta_j^{j-1}$.

Downstream links $m > j$. By the chain rule (A.26), $\gamma_{\text{new}}^{m-1} = \gamma^{j-1}\,(1-\theta_j^{j-1}-\varepsilon)\,\prod_{k=j+1}^{m-1}(1-\theta_k^{k-1})$. Property (iii) keeps the product factor unchanged; only the $(1-\theta_j^{j-1})$ factor changes, by a multiplicative factor $1-\varepsilon/(1-\theta_j^{j-1})$. Hence γ^{m-1} strictly decreases by approximately $\gamma^{m-1}\,\varepsilon/(1-\theta_j^{j-1})$. Since θ_m^{m-1} is unchanged, $L_m = \Gamma C_m/[\gamma^{m-1}\theta_m^{m-1}]$ strictly increases by approximately $L_m\,\varepsilon/(1-\theta_j^{j-1})$ for every $m > j$. In particular, L_b rises strictly above L^*.

Upstream links $m \leq j$ *(other than* j *itself).* For $m \leq j$, the joint default event $D_0 \cap \cdots \cap D_{m-1}$ does not involve agent j (since $m-1 < j$), and the conditional survival θ_m^{m-1} depends only on the joint distribution of $(m, 0, 1, \ldots, m-1)$, which is unchanged. Hence L_m is unchanged for $m < j$.

For ε small enough that L_j stays above L^*, the slack at j is reduced but not eliminated. L_b rises strictly above L^*, and other downstream link capacities also rise (staying above L^*). The new bottleneck level exceeds L^*. Strict improvement, contradicting σ's optimality.

Case 2: $b < j$ **(bottleneck upstream of the slack link).** Now perturb at the bottleneck b, not at j. (Perturbing at j would leave L_b unchanged, since γ^{b-1} does not involve agent j for $b \leq j$.) Substitute at b with $\theta_{b'}^{b-1} = \theta_b^{b-1} - \varepsilon$, $C_{b'} = C_b$, and θ_k^{k-1} unchanged for $k > b$.

Own link. L_b rises by $\approx L_b\,\varepsilon/\theta_b^{b-1}$, strictly above L^*.

Downstream links $m > b$. By (A.26), $\gamma_{\text{new}}^{m-1} = \gamma^{b-1}\,(1-\theta_b^{b-1}+\varepsilon)\,\prod_{k=b+1}^{m-1}(1-\theta_k^{k-1})$. Property (iii) keeps the product unchanged; the $(1-\theta_b^{b-1})$ factor increases by $\varepsilon/(1-\theta_b^{b-1})$ (multiplicatively). Hence γ^{m-1} strictly increases, and L_m strictly decreases, for every $m > b$ (including $m = j$).

Upstream links $m < b$. Unaffected.

For ε small enough, L_j drops slightly but remains above the new $L_b \approx L^*$ (since $L_j > L^*$ strictly by assumption). For other $m > b$ with $m \neq j$, the unique-bottleneck assumption gives $L_m > L^*$ strictly, so L_m stays above the new L_b for ε small.[4] The new bottleneck exceeds L^*. Strict improvement, contradicting σ's optimality.

[4]Specifically, ε must satisfy $\varepsilon < \min_{m>b}\,(L_m - L^*)/[\,L_m/(1-\theta_b^{b-1}) + L^*/\theta_b^{b-1}\,]$, where the minimum is over all $m > b$ on the chain. The right-hand side is strictly positive by the unique-bottleneck assumption ($L_m > L^*$ for every $m \neq b$), so a non-trivial ε always exists.

Combining the cases. At the optimum, no link can have $L_j > L^*$. Hence $L_j = L^*$ for every j on the chain, i.e., $C(k)/[\gamma^{i^*(k)-1}\theta_{i^*(k)}^{i^*(k)-1}] = L^*\Gamma^{-1}$ for every intermediary k. The equalization condition (12) holds. ■

Discussion: multi-bottleneck case. When several links are tied at the bottleneck level L^*, the single-position perturbations above are insufficient: in Case 2, lowering θ_b^{b-1} at one of the tied bottlenecks raises L_b but lowers L_m at every other tied bottleneck $m > b$, so the new bottleneck does not strictly improve. A combined perturbation at all tied bottleneck positions $b_1 < b_2 < \cdots < b_K$ resolves this. Let $\varepsilon_j > 0$ parameterize the magnitude of the perturbation $-d\theta_{b_j}^{b_j-1}$ at b_j, and let $M_{ij} \equiv \partial L_{b_i}/\partial\varepsilon_j$ denote the matrix of marginal effects on the tied bottlenecks. Under property (iii), the perturbation at b_j affects L_{b_i} for $i \geq j$ only via the $(1 - \theta_{b_j}^{b_j-1})$ factor in the chain rule (A.26), and it does not affect L_{b_i} for $i < j$ (since γ^{b_i-1} does not involve b_j). Hence M is *lower triangular* with strictly positive diagonal $M_{ii} = L_{b_i}/\theta_{b_i}^{b_i-1} > 0$, and is therefore automatically full-rank. A direction $(\varepsilon_1, \ldots, \varepsilon_K) > 0$ with $M\varepsilon > 0$ componentwise is constructed by routine forward substitution: pick $\varepsilon_1 > 0$, then for each $i = 2, \ldots, K$ pick ε_i satisfying $M_{ii}\varepsilon_i > \sum_{j<i} |M_{ij}|\varepsilon_j$. With each ε_i small enough that non-tied links remain above the new common bottleneck (the analog of the smallness bound footnoted in Case 2), the bottleneck level rises strictly, and the contradiction goes through as in the unique case.

A.10 Proof of Proposition 5

From comparing equations (15) and (16). ■

B Counterparty Risk under Explicit Rehypothecation

In the main text, the model does not distinguish between implicit and explicit collateral repledging. Under implicit repledging, collateral ownership transfers only upon the borrower's default, so the borrower faces no counterparty risk. Under explicit repledging, the pledged collateral is physically transferred along the chain at $t = 1$ when the borrowing takes place and returned at $t = 2$ upon debt repayment. This creates a counterparty risk: if the borrower's project succeeds but an intermediary on the chain is insolvent, the collateral may be locked in bankruptcy and not returned (Infante, 2019; Infante and Vardoulakis, 2021).

This appendix introduces counterparty risk and characterizes its impact on financing capacity and chain length. To evaluate the risk of collateral loss, the borrower needs to know the downstream chain, which stands in contrast to the baseline setting where such knowledge is not required and the borrower only needs to know the amount of credit offered by her immediate lender.

B.1 One Intermediary

We modify the setting of Section 3 as follows. Under explicit rehypothecation, the borrower transfers the pledged collateral to the intermediary at $t = 1$, who repledges it to the end lender. At $t = 2$, if the borrower repays, the collateral should flow back. However, if the intermediary is insolvent, the pledged collateral may be locked in the intermediary's bankruptcy process. Let $\delta \in [0, 1]$ denote the probability of collateral loss in this event.

At $t = 2$, four states arise:

(i) *Borrower succeeds, intermediary solvent* (probability $1 - 2q + q\phi$): The borrower repays. The intermediary passes the repayment to the end lender and returns the pledged collateral. The borrower's surplus relative to not borrowing is $(A - 1)L$ (investment gain, collateral recovered).

(ii) *Borrower succeeds, intermediary insolvent* (probability $q(1 - \phi)$): The borrower repays $R = L$, but the intermediary is insolvent. With probability $1 - \delta$, legal protections allow the collateral to flow back; with probability δ, the pledged collateral (expected value L) is lost. The borrower's expected surplus: $(A - 1)L - \delta L$.

(iii) *Borrower fails, intermediary solvent* (probability $q(1 - \phi)$): The borrower defaults. The intermediary seizes the pledged collateral. The borrower's surplus is $-L$ (collateral lost due to default on repayment, not counterparty risk). This is as in the baseline.

(iv) *Borrower fails, intermediary insolvent* (probability $q\phi$): Both default. The collateral passes to the end lender. The borrower's surplus is $-L$. This is as in the baseline.

We verify the state probabilities sum to one: $(1 - 2q + q\phi) + q(1 - \phi) + q(1 - \phi) + q\phi = 1 - 2q + q\phi + q - q\phi + q - q\phi + q\phi = 1$.

The borrower's expected surplus from intermediated financing under explicit rehypothecation is given by

$$\begin{aligned} S^{\delta} &= (1 - 2q + q\phi)\,(A-1)L \;+\; q(1-\phi)\,[(A-1)L - \delta L] \;+\; q(1-\phi)\,(-L) \;+\; q\phi\,(-L) \\ &= (A-1)L\,\big[(1-2q+q\phi) + q(1-\phi)\big] \;-\; q(1-\phi)\,\delta L \;-\; qL \\ &= (1-q)(A-1)L \;-\; qL \;-\; q(1-\phi)\delta L \\ &= \big[(1-q)(A-1) - q - q(1-\phi)\delta\big]\,L. \end{aligned} \tag{B.27}$$

The first two terms, $[(1-q)(A-1)-q]L$, are the baseline surplus from Lemma 2. The third term, $-q(1-\phi)\delta L$, is the expected counterparty-risk cost. Note that $q(1-\phi)$ is the joint probability that the borrower's project succeeds and the intermediary is insolvent: the probability of the borrower's success is $1-q$, and the intermediary's probability of being insolvent conditional on the borrower's success is $q(1-\phi)/(1-q)$ (derived by noting that the unconditional probability of the intermediary's failure is q, the joint failure probability is $q\phi$, and therefore the probability of intermediary failure conditional on borrower success is $(q-q\phi)/(1-q) = q(1-\phi)/(1-q)$; multiplying by $1-q$ gives $q(1-\phi)$).

Under direct financing (no intermediary), there is no counterparty risk and the surplus is $[(1-q)(A-1)-q]L_{\text{direct}}$, where $L_{\text{direct}} = \min\{\bar{U}, \Gamma C/q\}$. The chain dominates direct financing if and only if

$$\big[(1-q)(A-1) - q - q(1-\phi)\delta\big]\,L_{\text{chain}} \;>\; \big[(1-q)(A-1) - q\big]\,L_{\text{direct}}, \tag{B.28}$$

where $L_{\text{chain}} = \min\{\bar{U}, \Gamma C/\hat{q}\}$ is the intermediated credit capacity from (7). Rearranging:

$$\big[(1-q)(A-1) - q\big]\,(L_{\text{chain}} - L_{\text{direct}}) \;>\; q(1-\phi)\,\delta\,L_{\text{chain}}. \tag{B.29}$$

The left side is the benefit of intermediation (higher credit capacity). The right side is the counterparty-risk cost (increasing in δ, the collateral-loss probability, and in $1-\phi$, which governs how likely the intermediary is to fail when the borrower succeeds). When A is sufficiently large as stated in Assumption 1, the left side dominates for small δ, confirming that the chain remains beneficial when counterparty protections are adequate.

Discussion: the correlation trade-off. The parameter ϕ appears on *both* sides of the comparison (B.29). On the left side, the credit-capacity gain $L_{\text{chain}} - L_{\text{direct}}$ depends on ϕ through the bottleneck $\hat{q} = \max\{q\phi, q(1-\phi)\}$: ownership dilution is maximized at $\phi = 1/2$. On the right side, the counterparty-risk cost $q(1-\phi)\delta L_{\text{chain}}$ is decreasing in ϕ: when ϕ is high, the intermediary tends to fail *with* the borrower, so when the borrower succeeds, the intermediary is likely solvent and the collateral returns safely. In the baseline ($\delta = 0$), the optimum is $\phi = 1/2$. Under counterparty risk ($\delta > 0$), the optimum shifts toward $\phi > 1/2$—that is, toward intermediaries whose solvency is more positively correlated with the borrower's—because higher correlation buys collateral safety at the cost of a tighter downstream bottleneck.

B.2 Multiple Intermediaries

We extend the analysis to a chain with $n-1$ intermediaries under explicit rehypothecation and the conditionally independent survival of intermediaries as in the setting of Section 4.3.[5] We maintain the symmetric assumption that each intermediary's unconditional failure probability equals q, as in the two-firm setting of Section 3.

The pledged collateral is rehypothecated along the entire chain: borrower $\to$ intermediary 1 $\to \cdots \to$ end lender. When the borrower repays at $t = 2$, the repayment passes through the chain (each intermediary passes it downstream to repay her own lender), and the collateral flows back (end lender $\to \cdots \to$ intermediary 1 $\to$ borrower). If any intermediary on the chain is insolvent, she may be unable to pass the repayment downstream or coordinate the collateral return—the chain of actions breaks at that point and the borrower may not recover the collateral.[6]

[5] We adopt the conditionally independent survival of intermediaries for clean closed-form expressions. The qualitative conclusions of this subsection—that counterparty risk weakly shortens the optimal chain and that the joint failure correlation structure mediates the counterparty-risk cost—extend to general joint failure distributions, though the closed-form expression for π_n does not.

[6] We assume δ is the same regardless of chain length. In practice, δ may increase with chain length because longer chains involve more complex unwind procedures.

Each intermediary k's probability of failing, conditional on the borrower's success, is

$$P(k \text{ fails} \mid 0 \text{ succeeds}) = \frac{q\,\theta(k)}{1-q}, \tag{B.30}$$

where $\theta(k) = 1 - \phi(k)$ is k's survival probability conditional on the end borrower 0's default. This expression is derived from the law of total probability: $P(k \text{ fails}) = P(k \text{ fails} \mid 0 \text{ fails})\, q + P(k \text{ fails} \mid 0 \text{ succeeds})\,(1-q)$. Using $P(k \text{ fails}) = q$ under homogeneous marginal failure probability for all firms (including the end borrower and intermediaries) and $P(k \text{ fails} \mid 0 \text{ fails}) = 1 - \theta(k)$, we solve: $P(k \text{ fails} \mid 0 \text{ succeeds}) = [q - q(1-\theta(k))]/(1-q) = q\theta(k)/(1-q)$.[7]

The expression (B.30) is increasing in $\theta(k)$: an intermediary with high $\theta(k)$ (high survival conditional on the end borrower's failure, i.e., low correlation with the borrower) is more likely to fail independently when the borrower succeeds, because its failure is less tied to the borrower's.

Under the weakest-link model—the chain breaks if any intermediary is insolvent—the probability that the collateral is at risk, conditional on the borrower's success, is

$$\pi_n \;\equiv\; 1 - \prod_{k=1}^{n-1}\left[1 - \frac{q\,\theta(k)}{1-q}\right], \tag{B.31}$$

which is strictly increasing in the number of intermediaries.

The borrower's modified surplus with $n-1$ intermediaries under explicit rehypothecation is

$$S_n^{\delta} \;=\; \big[(1-q)(A-1)-q\big]\,L(n) \;-\; (1-q)\,\pi_n\,\delta\,L(n) \;=\; \big[(1-q)(A-1)-q-(1-q)\pi_n\delta\big]\,L(n), \tag{B.32}$$

where $L(n)$ is the credit capacity of the n-link chain from (9) with $n-1$ intermediate lenders and one end lender, and $(1-q)\pi_n$ is the unconditional probability that the borrower succeeds and at least one intermediary is insolvent. The first term is the baseline benefit; the second is the expected cost of counterparty risk, which grows with n, the number of links of the chain.

For the chain with $n-1$ intermediaries to dominate one with $n-2$ intermediaries, the marginal benefit of adding the $(n-1)$-th intermediary must exceed the marginal counterparty-risk cost, i.e., $S_n^{\delta} > S_{n-1}^{\delta}$, or equivalently

$$\big[(1-q)(A-1)-q-(1-q)\pi_n\delta\big]\,L(n) \;>\; \big[(1-q)(A-1)-q-(1-q)\pi_{n-1}\delta\big]\,L(n-1). \tag{B.33}$$

[7] When intermediaries have heterogeneous marginal failure probabilities $q_k \neq q$, (B.30) becomes $P(k \text{ fails} \mid 0 \text{ succeeds}) = [q_k - q_k(1-\theta(k))]/(1-q) = q_k\theta(k)/(1-q)$. The analysis extends with q replaced by q_k.

In the baseline model ($\delta = 0$), this reduces to $L(n) > L(n-1)$: the chain extends whenever credit capacity strictly increases. Under counterparty risk ($\delta > 0$), the condition is strictly harder to satisfy because the left side is penalized by the higher π_n and the right side is less penalized by the lower π_{n-1}. Formally, rearrange (B.33):

$$\big[(1-q)(A-1)-q\big]\,\big[L(n)-L(n-1)\big] \;>\; (1-q)\,\delta\,\big[\pi_n\,L(n)-\pi_{n-1}\,L(n-1)\big], \qquad \text{(B.34)}$$

where $(1-q)(A-1)-q > 0$. The left side is the marginal ownership-dilution benefit (the same term that drives the baseline model). The right side is the marginal counterparty-risk cost, which is strictly positive because both $\pi_n > \pi_{n-1}$ and $L(n) \geq L(n-1)$.

Since the right side of (B.34) is positive and absent in the baseline, the optimal chain under counterparty risk is *weakly shorter* than in the baseline model: some intermediaries whose addition would have been profitable in the baseline are no longer profitable once the marginal counterparty-risk cost is taken into account. The chain terminates where the marginal ownership-dilution benefit no longer exceeds the marginal counterparty-risk cost.

Beyond chain length, this extension yields another natural empirical prediction: in environments with weaker collateral protection, credit chains should rely on intermediaries whose solvency is more positively correlated with that of the end borrower so that when the end borrower survives the intermediaries are likely to survive as well and the collateral will not be lost.

B.3 Real-World Mitigation of Counterparty Risk

Several institutional features reduce δ toward zero, bringing the setting closer to the baseline model:

- *Bankruptcy-remote SPVs.* In structured-finance transactions, collateral is held within bankruptcy-remote special-purpose vehicles that legally isolate it from the intermediary's estate, ensuring the collateral is returned even if the intermediary is insolvent.

- *Tri-party custody.* In tri-party repo markets, a custodian bank holds the collateral on behalf of both the cash borrower and the cash lender. This custodial arrangement reduces bilateral counterparty exposure and facilitates orderly unwinds when an intermediary fails.

- *Repo safe-harbor provisions.* In many jurisdictions, repurchase agreements are exempt from the automatic stay in bankruptcy, allowing the non-defaulting party to seize or return collateral immediately without waiting for bankruptcy proceedings.

- *Segregation requirements.* Regulations such as margin segregation rules for derivatives and customer-protection rules for broker-dealers require that posted collateral be held in segregated accounts, separate from the intermediary's proprietary assets.

When A is sufficiently large and δ small, the counterparty-risk cost in (B.27) and (B.32) becomes negligible relative to the benefit of ownership dilution from chain extension, and the predictions of the baseline model are approximately correct. The baseline model applies most directly to settings where legal protections are strong (as in tri-party repo) and to implicit repledging (our default interpretation of the model), where $\delta = 0$ by construction. We want to highlight that when δ is significant, this extension of our model with counterparty risk generates a new, testable prediction: intermediation chains should be shorter and involve more highly correlated intermediaries in settings with weaker protections against counterparty risk.