

Formation of Optimal Interbank Networks under Liquidity Shocks

Daniel E. Rigobon* Ronnie Sircar†

October 1, 2025

Abstract

We study the formation of an optimal interbank network in a model where banks control both their supply of liquidity, through cash holdings, and their exposures to other banks' risky projects. The value of each bank's project may suddenly decline due to liquidity shocks, which are impactful if the bank has insufficient cash holdings when they occur. In two distinct decentralized and centralized settings, we solve the system-wide optimal control problem and obtain explicit formulas for the unique optimal allocations of capital. Both of the resulting networks exhibit a 'core-periphery' structure in equilibrium. However, in the decentralized setting, banks elect to hold less cash reserves compared to the centralized setting, leading to greater susceptibility to liquidity shortages. We characterize the behavior of the planner's optimal allocation as the size of the system grows. Surprisingly, the price of anarchy is of constant order. Finally, we derive co-investment requirements that allow the decentralized system to achieve the planner's optimal level of risk. In doing so, we find that banks in the network's core are subjected to the highest co-investment requirements – ensuring that they are sufficiently incentivized to hold significant amounts of cash. Our analysis may inform regulators' requirements on banks' liquidity reserves, as have been debated in the wake of the 2023 regional banking crisis in the US.

Keywords: Banking networks, stochastic intensity control, centralized and decentralized equilibria

*Department of Operations Research & Financial Engineering (ORFE), Princeton University.
drigobon@princeton.edu

†Department of Operations Research & Financial Engineering (ORFE), Princeton University.
sircar@princeton.edu

1 Introduction

In this paper, we construct a tractable and flexible continuous-time model of a financial system in which banks are exposed to both their own and possibly each others’ liquidity shocks. We analyze and obtain closed-form solutions for this model in two distinct settings: market equilibrium (decentralized) and the social planner’s problem (centralized). A further contribution is being able to quantify, again in closed form, the extent of inefficiencies associated with market equilibrium outcomes through the characterization of the ‘Price of Anarchy’. In contrast to much of the existing literature on banking networks, which assumes that the linkages between financial institutions are exogenous, we allow this network to form endogenously based on optional decisions taken by participating institutions.

The model proceeds as follows. Banks may specialize in different activities; some collect a large number of deposits, whereas others specialize in revenue generation. This heterogeneity is modeled by unique, proprietary, investment opportunities (e.g. a portfolio of commercial loans) available to each bank. We assume that these opportunities are scalable, and accessible to other banks within the system. For example, a deposit-collecting bank i can obtain the large returns of investment bank j ’s unique revenue-generating operations by directly investing into project j . We note that this construction is similar to the single-period models of Rochet and Tirole (1996), Acemoglu et al. (2015b), and Erol and Ordonez (2017), and we also refer to these unique investment opportunities as ‘projects’. In these previous models and ours, the riskiness of these projects is tied to decisions taken by the associated bank.

Although these projects may accrue large rates of return, they are subject to a degree of risk. More precisely, a bank’s project may be impacted by unpredictable liquidity shocks of random magnitude. These shocks are assumed to represent, for example, additional liquidity required for a bank’s project to avoid losses, such as occurs in Rochet and Tirole (1996) and Acemoglu et al. (2015b), or the withdrawal of a large number of deposits, which triggered the decline and ultimate failure of Silicon Valley Bank in 2023. If the size of such a shock exceeds the bank’s current cash holdings, then the project’s value instantaneously drops. We refer to such an event as a liquidity shortage. When this occurs, all banks investing in this project suffer losses proportional to their stake in it. Therefore, conditioned on the arrival of a liquidity shock, a bank’s cash holdings determines their project’s level of risk. Conversely, if a bank has sufficient cash on hand when a liquidity shock arrives, their project suffers no losses. Figure 1.1 illustrates the relationship between a bank’s cash holdings, the distribution of a liquidity shock’s size, and the probability of liquidity shortage conditioned on shock arrival.




Figure 1.1: For a single bank, the relationship between the cumulative distribution function (CDF) of the size of a liquidity shock, their cash holdings, and the conditional probabilities of (in)sufficient liquidity.

In 2023, there was some notable failures of mid-size US banks, particularly Silicon Valley Bank (SVB), First Republic Bank (FRB) and Signature Bank. These seem to have been initiated by interest rate increases by the Federal Reserve in its efforts to reduce inflation, leading to losses in value to these banks' bond portfolios. Rising interest rates also led to higher financing costs on many of their large institutional clients within the Silicon Valley tech sector, who then began withdrawing their deposits to meet their own liquidity needs. This impact on banks was then exacerbated, as other large depositors became concerned about the banks going into distress and their inability to recover beyond the \$250K insured by the FDIC. They withdrew large amounts very quickly, typically using online banking apps, and often outside of business hours and over weekends. On March 8, 2023, depositors withdrew \$42 billion from SVB in one day, and two days later Federal regulators seized control of the bank.

However, these banks did not have sufficient cash holdings to meet depositors' demand, with significant amounts of capital tied up in long-term treasury bonds. These bonds had been purchased during an earlier low interest rate era and were liquidated at considerable loss. Ultimately this was not enough to restore depositor confidence, and the banks ended up insolvent due to lack of liquidity. The run on SVB likely lowered confidence in similar tech industry concentrated banks such as FRB, which led to a run on these banks. While our model was not designed to capture this kind of contagion effect, it does indeed focus on the impact and prevention of liquidity shocks through cash holdings in an inter-connected banking network.

A key focus of this paper is that each bank endogenously chooses to allocate its capital between cash holdings and other banks' risky projects. To that end, we will study the optimal capital allocations for two extreme settings of the financial system. First, we consider the *decentralized* case – wherein each bank freely allocates their capital with pure self-interest. They seek only to maximize their utility of wealth at a given terminal time. We note that this setting reflects a market equilibrium. Second, we consider a *centralized* setting – where a single social planner makes the allocation decisions for all banks concurrently and aims to maximize the sum of individual banks' utilities.

In both decentralized and centralized cases, we derive the dynamic programming equations for the respective value functions, and explicitly compute the optimal allocations, which yield a core-periphery network structure in both cases. Under stricter technical conditions, we also prove uniqueness. We observe a discrepancy between the optimal allocations computed in both settings: the social planner chooses to hold a greater amount of cash. This occurs because our model captures a simple negative externality: when individualistically determining how much cash to hold, a bank sets the level of risk experienced by external investors in its project. An individual bank, operating in its own interests, fails to internalize these external investors' exposures when choosing their amount of cash to hold. In contrast, the planner is cognizant of the network-wide welfare and acts accordingly by reducing the level of risk through increased cash holdings according to each project's degree of external investment. By design, the social planner achieves the welfare-maximizing (i.e. first-best) allocation for the financial system. As a consequence of this discrepancy, liquidity shortages are less likely to occur under the social planner than under decentralized behavior. However, we also observe that interbank exposures are larger in the centralized setting, and hence each liquidity shortage becomes more damaging to the system. This tension between the likelihood and severity of extreme events bears a resemblance to the 'robust yet fragile' observation made by Gai and Kapadia (2010) and Acemoglu et al. (2015a). In particular, we find that this feature is associated with the socially optimal allocation of capital in the financial system.

We also study how the two optimal allocations differ as the financial system's size increases. Two natural points of comparison are: 1) the difference in, and 2) the ratio of, social welfare between both settings. The former comparison measures the nominal size of the inefficiency, and

the latter its relative size (which we refer to as the ‘price of anarchy’, as in Papadimitriou (2001)). Perhaps counter-intuitively, we find that the price of anarchy remains bounded by a constant as the size of the system grows. Namely, the nominal size of the system’s inefficiency grows at the same rate as the social welfare itself. These results are first derived theoretically, and also verified in simulations. Finally, we show that it is possible to alter banks’ investment constraints to replicate the social planner’s optimal allocation, which requires increasing restrictions only on banks in the network’s core.

There are several interesting takeaways from this paper. First, we are able to explicitly solve a continuous-time network optimal control problem and characterize the loss in efficiency due to decentralized behavior. Our model contains: i) multiple interacting agents, each solving their own utility maximization Merton problem, ii) uncertainty driven by jump processes rather than standard Brownian motion, and iii) agents’ ability to control the intensity of their jump process. To the best of our knowledge, explicit solutions in a model with multiple controlled stochastic intensities is novel. Second, we find that the planner’s optimal allocation leads to low-frequency and high-magnitude losses in the system. This may imply that the ‘robust-yet-fragile’ feature of financial networks is, in some sense, socially optimal. However, we see that the planner compensates for the larger-magnitude losses by ensuring they are less likely. As a result, the planners’ solution is associated with larger interbank investments throughout the system. Additionally, in both settings we see that the endogenous financial networks exhibit a ‘core-periphery’ structure, where only a subset of banks’ projects are invested in by other banks. Intuitively, we show that only banks in the network’s core must be more strictly regulated to replicate the planner’s optimal allocation. This suggests there is greater value in regulation of banks in the endogenously-determined core of the network, over those in the periphery.

This paper is organized as follows. Section 1.1 reviews several related branches of literature and background. Section 2 introduces the model of the financial system, the dynamics of each instrument, and the control problem posed to each bank. In the first part of our main results, Section 3 derives the endogenous equilibria for interbank networks in the decentralized (Section 3.2) and centralized (Section 3.3) settings. We compare these two optimal allocations in Section 4, including an asymptotic analysis of the price of anarchy. Finally, Section 5 concludes with a discussion of our results and directions for future work.

1.1 Related Literature

A strong motivation for this paper follows from the systemic risk literature, where much of the existing work assumes a given or exogenous network structure for the financial system. An early paper by Allen and Gale (2000) studies several stylized structures of interbank claims, and finds that the structure determines whether or not a local liquidity shock propagates throughout the system. Later papers seek to answer similar questions with distinct models. For instance Gai and Kapadia (2010) and Gai et al. (2011) find that systemic liquidity crises can emerge in highly interconnected financial networks, albeit with low probability. Caccioli et al. (2014) present a model in which firms’ overlapping portfolios can lead a single default to cause mark-to-market losses throughout the system – perhaps leading to additional defaults. In Elliott et al. (2014), firms directly own claims on each others’ assets and suffer sudden bankruptcy losses if their valuation falls below a threshold. Battiston et al. (2012) studies a continuous-time process representing financial robustness, and allows its evolution to depend on a given financial network. Finally, several papers including Amini and Minca (2016), Detering et al. (2019, 2020) and Detering et al. (2021) seek to characterize the asymptotic behavior of contagion cascades in random inhomogeneous networks as the system’s size grows. In their respective studies, these different mechanistic models are investigated both

theoretically and in simulations. In addition, the two-period model of Ramírez (2022) shows that under extreme levels of uncertainty, banks fail to internalize their effect on contagion cascades. However, the explicit or implicit networks in these papers share one common feature – they are fixed or generated according to canonical random graph models. As previously highlighted, we believe this assumption may not be realistic; institutions in the financial system make optimal investment decisions, and the resulting network is endogenous – not random or otherwise pre-specified. In contrast to this branch of the literature, our model enables us to investigate how the organization and fundamental parameters of the financial system can lead to the emergence and scale of its inefficiencies.

We use several tools from continuous-time portfolio optimization in this work, which begins with the foundational papers of Merton (1969, 1971). Merton studies the optimal portfolio allocation between risk-free and risky assets for a investor who maximizes their expected discounted utility of consumption. In these models, the returns of each risky asset are driven by correlated Brownian motions. Following from Merton’s seminal papers, there is a wealth of literature on extensions of the original problem; see Rogers (2013) and references therein. The stochastic control techniques we use in this paper for deriving the optimal allocations will be similar to Merton’s original work and its subsequent branch of literature, but here we will be studying a financial system in which *all* participants are simultaneously determining their optimal allocations of wealth – not only an individual. Moreover, to the best of our knowledge, the ability for multiple players to control the jump intensity of their the risk to their asset’s returns has not been previously studied in the area of portfolio optimization. A related mean field game with a continuum of cryptocurrency miners controlling the intensity of their discovery arrival process is analyzed in Li et al. (2024).

There are, however, several papers that study an individual who incurs a cost to control the intensity of a jump process, such as Biais et al. (2010), Pagès and Possamai (2014), Capponi and Frei (2015), Hernández Santibáñez et al. (2020), and Bensalem et al. (2020). These studies focus on Principal-Agent models and largely analyze the optimal contract and behavior. Moreover, they focus on the presence of moral hazard, where the Principal is unable to observe the Agent’s efforts. Our mathematical approach for determining a bank’s optimal cash holdings is similar to the models used in these papers. However, there are a few important differences. First, we study these optimizations performed simultaneously within a large system, and second, we focus on the inefficiencies that arise when banks optimize only in their self-interest. Additionally, our setting assumes perfect information, so there is no moral hazard.

We note that the high-level ideas in this paper are similar to the literature on optimal network formation. For early work in this area, see Jackson and Wolinsky (1996) and Bala and Goyal (2000), where the authors present a process by which individuals choose to create edges with each other in a game-theoretic model. In these studies, individuals must balance a trade-off between the cost of forming an edge and the rewards associated with the edge. Our paper differs primarily from these studies through our emphasis on the financial features of the model, and the fact that edges are cost-less to form. A more realistic extension of our work would certainly incorporate these kinds of fixed costs.

Most closely related to this paper is the study of endogenous financial networks, including Erol and Ordonez (2017), Zawadowski (2013), Bluhm et al. (2014), Acemoglu et al. (2015b, 2021), Babus (2016) and Farboodi (2021). Erol and Ordonez (2017) studies the formation of linkages along which liquidity can flow between banks in a system, where each bank has access to a proprietary project. They find that excessive restrictions on banks’ liquidity reserves increase the level of systemic risk. Our model primarily differs from this work in that linkages between banks do not provide liquidity to either counterparty, and there is no mechanism for contagion. The work of Zawadowski (2013) shows that individual banks may fail to achieve the socially-optimal outcome by not buying

insurance against their counterparties' default. While the author's model differs greatly from ours, we similarly find that individual banks' optimal behavior fails to take into account an externality on the system. A model by Babus (2016) presents an extension of Allen and Gale (2000). Her model allows banks to make optimal lending and borrowing decisions to redistribute liquidity throughout the system, and a highly-connected network is again found to be the most resilient to contagion. We share the idea of idiosyncratic liquidity shocks, but also study the planner's optimal allocation and compare it to the case where banks make individually optimal decisions.

The papers most similar to our own are Bluhm et al. (2014), Acemoglu et al. (2015b, 2021) and Farboodi (2021). Our model is fundamentally different from those in these studies – which are either static or consist of three distinct time periods. In contrast, we analyze the endogenous financial network problem in a continuous-time environment, which allows us to leverage the powerful mathematical tools of stochastic optimal control. First, Bluhm et al. (2014) construct a model of optimal interbank lending where banks face both liquidity and capitalization requirement constraints. In their model, both the interbank lending amounts and the market prices are determined endogenously. The authors show that contagion can occur (1) directly as a result of counterparty losses in the event of a default, or (2) indirectly through the mark-to-market losses incurred by a bank's portfolio in the event of a fire sale. The authors largely focus on numerical and simulation results, while we are able to provide explicit formulas for the equilibria we study. Moreover, our model endogenizes the initial sources of disruption to the system.

The contribution of Farboodi (2021) characterizes how banks optimally lend to each other within a financial system where there is a strong incentive to serve as intermediaries within the chain of lending. In her model, as in ours, an interbank exposure will also allow the originating bank to access the surplus generated by a risky investment of an associated bank. She shows that the resulting network can have a core-periphery structure, and that due to the benefit of intermediation, banks' private incentives can fail to achieve the socially optimal outcome. Although there are similarities between this paper and ours, we do not focus on the incentive of intermediation, but instead on banks' optimal decisions to reduce the riskiness of their investments. Our results also replicate the core-periphery feature in her paper – a small subset of banks with highly profitable investment opportunities form the financial network's core.

Finally, Acemoglu et al. (2015b, 2021) endogenize both the decision of interbank lending and also the interbank interest rates. Acemoglu et al. (2021) build a model whereby linkages between banks provide liquidity throughout the system, and the anticipation of contagion leads to scarce or non-existent credit. They study the relationship between the distribution of liquidity shocks and the occurrence of systemic credit freezes, along with the role of different policy interventions in curbing these inefficiencies. In contrast to this work, our paper focuses on the investment decisions of participating banks in the financial system, where their liquidity shocks must be addressed by their own cash holdings. Acemoglu et al. (2015b) mirrors the work of Rochet and Tirole (1996), where banks exchange deposits to finance a project that yields high rate of return if run to conclusion, or low returns if liquidated prematurely. A bank faces external liabilities that may require them to liquidate these projects – thereby passing losses onto its creditors. The authors find that the optimal contracts do indeed consider the first-order network effects, wherein a risk-taking bank must pay large interest rates to its creditors. However, these do not account for the 'financial network externality', which can negatively affect banks that are not party to the contract. It follows that the resulting financial network may not be efficient (i.e. welfare-maximizing). While their models of interbank dynamics are similar to ours, the authors' analysis is largely focused on stylized networks in which equilibria are shown to exist. In this paper, we will instead allow the sparsity structure of the financial network to be endogenously determined by the interbank investment opportunities and other banks' decisions.

2 Model

Consider a financial system consisting of n different banks. Let $(\Omega, \mathcal{E}, \mathbb{P})$ be a probability space, containing n mutually independent Poisson processes $\tilde{N}_t^1, \dots, \tilde{N}_t^n$, $t \geq 0$, each of which has corresponding intensity $\theta_1, \dots, \theta_n > 0$. These counting processes will be used to indicate the arrival times of liquidity shocks to each respective bank. Define \mathcal{F} to be the filtration generated by the full set of jump processes. Hence, we obtain the filtered probability space $(\Omega, \mathcal{E}, \mathcal{F}, \mathbb{P})$. The net capitalization (i.e. net value or wealth) of bank i is given by the non-negative stochastic process $\{X_t^i\}_{t \geq 0}$, whose dynamics are described in the following.

2.1 Investment Opportunities

First, the financial system contains a long-term bond, which accumulates a constant, fixed rate of return $r \geq 0$ over all time. Therefore its value at time t , denoted S_t^0 , evolves according to the ordinary differential equation $\frac{dS_t^0}{S_t^0} = r dt$. The rate of return r will be referred to in this paper as the ‘risk-free rate’, since it is unaffected by any stochastic factors. However, we note that holdings in the bond are not free from *liquidity* risk. Namely, we will assume that the bond cannot be used to satisfy short-term liquidity needs. We think of the rate r as capturing the rate of return for long-term and less liquid fixed income instruments, such as 5- or 10-year Treasury bills. Appendix C.2 contains an extension of our model where this assumption is weakened and the bond can be liquidated for a fraction of its face value. Computationally, we see no significant differences in large enough banking networks.

Each bank i has access to a unique set of investments henceforth referred to as a ‘project’, e.g. a collection of commercial loans. These projects are made available for investment to all other banks $j \neq i$ in the system. However, these projects may not be shorted, and any capital invested in them cannot be rapidly liquidated. More precisely, no bank can use this capital to meet any short-term liquidity needs. While these projects accumulate large constant rates of return for investors, they will incur losses when the associated bank suffers a liquidity shortage. If such an event occurs, then the project’s value immediately drops to some non-zero fraction of its prior value. For instance, it is plausible that a bank’s revenue-generating operations intermittently require additional liquidity to cover a position or meet regulatory requirements. Failure to do so may lead to the bank’s inability to realize the investment’s gains, or may even directly cause losses.

Let S_t^i denote the time- t value of a single unit of capital invested in bank i ’s project. Its dynamics are given by

$$\frac{dS_t^i}{S_t^i} = (\mu_i + r) dt - \phi_i dN_t^i, \quad i = 1, \dots, n. \quad (2.1)$$

We assume that $\mu_i > 0$, so this project has deterministic rate of return greater than r . The jump process N_t^i is obtained by performing a thinning of the shock arrival process \tilde{N}_t^i , and is described in the next paragraphs. The increment dN_t^i takes on values in $\{0, 1\}$, and is equal to one if and only if bank i experiences a liquidity shortage at time t . Finally, $0 < \phi_i < 1$ represents the fraction of project value lost when such an event occurs.

All banks in the system may experience liquidity shocks; if sufficiently large, these shocks adversely affect the value of their project. A key feature of this paper is each bank’s ability to control their project’s susceptibility to such events – by holding a greater supply of liquidity in the form of cash holdings. In our model, this is represented through bank i ’s ability to influence the intensity of the jump process N_t^i that appears in (2.1).

A bank may hold a non-negative amount of their capital as in cash which does not accrue any interest. However, these cash holdings are the *only* source of short-term liquidity within the system,

and are the sole tool by which a bank can hedge against the arrival of liquidity shocks. Namely, if a liquidity shock exceeds bank i 's cash holdings, their project experiences a substantial loss in value, which impacts both their own wealth and that of other banks investing in i 's project. The jump increment dN_t^i in (2.1) equaling one represents the arrival of a shock that overwhelms bank i 's supply of cash, and its construction follows from a probabilistic model of liquidity shocks and a bank's cash holdings.

Recall that our filtered probability space contains n independent time-homogeneous Poisson processes \tilde{N}_t^i , with rates $\theta_i > 0$. At time t , if \tilde{N}_t^i jumps, then bank i experiences a liquidity shock of size $\zeta_t^i \cdot X_t^i$, where the random variable ζ_t^i is \mathcal{F}_t -measurable. We assume that these shocks are proportional to a bank's wealth, and each ζ_t^i is independently and identically distributed according to the cumulative distribution function (CDF) $F_i(\cdot)$. For convenience of notation, the complementary CDF of ζ_t^i is denoted by $\bar{F}_i(\cdot) = 1 - F_i(\cdot)$.

Let $c_t^i \geq 0$ denote the fraction of bank i 's capital held as cash holdings at time t . When the shock to bank i is larger than these reserves (i.e. $\zeta_t^i > c_t^i$), they undergo a liquidity shortage and all investors in their project i suffer an instantaneous return of $-\phi_i$ on their investment. In particular, if $c_t^i = 0$, then any liquidity shock to bank i at time t , no matter how small, results in losses.

The jump process N_t^i is constructed by independently flipping a coin at every jump arrival time of \tilde{N}_t^i , with failure probability given by $p_t^i = \bar{F}_i(c_t^i)$. Observe that $p_t^i = \mathbb{P}\left(\zeta_t^i > c_t^i \mid d\tilde{N}_t^i = 1\right)$. We let $dN_t^i = 1$ if and only if the flip is a failure. It follows from the thinning properties of Poisson processes that the instantaneous rate at time t of the process N_t^i is equal to $\theta_i \bar{F}_i(c_t^i)$.¹ The second component of the rate, $\bar{F}_i(c_t^i) = \mathbb{P}(\zeta_t^i > c_t^i)$, is the probability that bank i suffers a liquidity shortage, conditioned on the time- t arrival of a liquidity shock with CDF $F_i(\cdot)$. See Figure 1.1 for an illustration.

Remark 2.1. *In the main text, we assume that liquidity shocks smaller than the bank's cash holdings (i.e. when $\zeta_t^i \leq c_t^i$ and $d\tilde{N}_t^i = 1$) do not affect the bank's overall wealth. Appendix C.1 relaxes this assumption, and this extended model is still solvable up to first-order conditions, which must be solved numerically. As long as the expected shock size is small, there is no significant difference caused by this relaxation.*

Finally, we will require a few technical conditions on F_i :

Assumption 1. *We assume that each F_i is absolutely continuous with respect to the Lebesgue measure. Its density is given by $f_i(\cdot) = F_i'(\cdot)$, which is assumed to be fully supported on \mathbb{R}_+ , and monotonically decreasing (i.e. $f_i'(\cdot) < 0$).*

If $f_i(\cdot)$ had compact support, then it would be possible for a bank's project to be riskless with large enough cash holdings. Since the return of this project would be greater than the risk-free rate, this would lead to all other banks profiting infinitely by borrowing at the risk-free rate and investing in this riskless project. While the problem may remain analytically tractable, this outcome is not of practical interest. Our assumption that the density is monotonically decreasing will be used to establish uniqueness of the optimal controls.

2.2 Dynamics of Wealth

In this model, each bank may invest in any other banks' projects. Let $w_t^{ji} \geq 0$ denote the fraction of bank j 's capital invested in bank $i \neq j$'s project. The return on this investment is described by (2.1).

¹This result is a consequence of the thinning properties of Poisson processes. See, for instance, Theorem 1 in Lewis and Shedler (1979). If $\{c_t^i\}_{t \geq 0}$ is adapted (as we will require), then conditioned on time t , the previous jump process $\{N_s^i\}_{s \in [0, t]}$ has the desired rate function.

Recall that c_t^i equals the fraction of bank i 's wealth held as cash holdings, which accumulates no return over time.

The final component influencing bank i 's wealth is their degree of investment in their own project, which is equal to a given and fixed fraction of their current wealth. Unlike exposures to other banks' projects, we assume that this quantity cannot be optimized by the bank. For instance, it may be the case that bank i is required by a regulator to be an investor in its own project. In principle, we could imagine allowing bank i to also control their exposure to their own project. However, doing so introduces the possibility for multiple equilibria, as shown in Appendix C.3.

We will use $\frac{\eta_i}{\phi_i}$ to denote the fraction of bank i 's wealth that is invested in their own project. This implies that bank i loses a constant fraction η_i of its total wealth whenever they suffer a liquidity shortage. The parameter η_i captures the severity of losses on the distressed bank – in the extreme case of $\eta_i = 1$, a single liquidity shortage will wipe out bank i . Conversely, if $\eta_i = 0$, then bank i has no stake in their project and is wholly unaffected by its dynamics. We will take $0 < \eta_i < 1$, away from the two extreme cases.

Due to self-financing constraints, the remaining $X_t^i(1 - c_t^i - \sum_{j \neq i} w_t^{ij} - \frac{\eta_i}{\phi_i})$ units of wealth are held as long or short positions in the interest-bearing bond. Putting together the dynamics for each component of bank i 's wealth, we see that X_t^i , follows

$$\frac{dX_t^i}{X_t^i} = \left(1 - c_t^i - \sum_{j \neq i} w_t^{ij} - \frac{\eta_i}{\phi_i} \right) \frac{dS_t^0}{S_t^0} + \sum_{j \neq i} w_t^{ij} \frac{dS_t^j}{S_t^j} + \frac{\eta_i}{\phi_i} \frac{dS_t^i}{S_t^i}, \quad i = 1, \dots, n.$$

Using (2.1) and the dynamics of S_t^0 , we obtain the following simplified expression:

$$\frac{dX_t^i}{X_t^i} = \left((1 - c_t^i)r + \sum_{j \neq i} w_t^{ij} \mu_j + \frac{\eta_i}{\phi_i} \mu_i \right) dt - \sum_{j \neq i} w_t^{ij} \phi_j dN_t^j - \eta_i dN_t^i, \quad i = 1, \dots, n. \quad (2.2)$$

A novel contribution of this paper is the control c_t^i ; while there is no return accumulated by cash holdings, they serve to reduce the likelihood that bank i suffers a liquidity shortage, which would cause them to lose a fraction η_i of their wealth.

We say that (c^i, w^i) is in $\mathcal{A}_{s,t}^i$, the set of admissible controls for bank i between times s and t , if it is adapted to the filtration \mathcal{F} and satisfies both $c_u^i \in \mathbb{R}_+$ and $w_u^{ij} \in [0, \phi_j^{-1})$ for all $u \in [s, t]$ and $j \neq i$. The upper bound on w_u^{ij} ensures that wealth will always remains positive.

In the first decentralized setting we study, all banks seek to maximize, over their cash holdings and investments in other banks' projects, their own expected utility of wealth at a common terminal time $T < \infty$: equal to $\mathbb{E}[U_i(X_T^i)]$. As is relatively standard in the literature, a bank's utility function $U_i \in \mathcal{C}^\infty(\mathbb{R}_+)$ is assumed to have constant relative risk aversion of $\gamma_i > 0$, which yields the following family of utility functions:

$$U_i(x) = \begin{cases} \frac{x^{1-\gamma_i}}{1-\gamma_i} & \gamma_i > 0, \gamma_i \neq 1 \\ \log x & \gamma_i = 1. \end{cases} \quad (2.3)$$

For generality, we will allow banks to have heterogeneous risk-aversion coefficients γ_i in the decentralized network formation of Section 3.2, while in Section 3.3 a central planner will require all banks to have logarithmic utility.

3 Decentralized and Centralized Financial Networks

We consider two distinct organizations of the financial system. In the first, banks operate only in their self-interest – seeking to maximize their own expected terminal utility. We call this the *decentralized* setting, as there is no coordination between banks. Instead, each bank’s optimal allocation reflects their best response to the others’ decisions. On the other hand, the *centralized* setting in Section 3.3 will consider the perspective of a single social planner who determines all banks’ allocations to maximize welfare – as measured by the sum of all banks’ utilities.

Both allocations are important to consider. The decentralized optimum reflects a game-theoretic or market-mediated equilibrium of the financial system, where each bank chooses their controls optimally given all others’ actions. Therefore, from the perspective of individual banks this is a stable allocation. In contrast, the centralized optimum reflects the maximum total welfare that could exist in the financial system if banks’ capital allocations were coordinated by a central planner. We will study the differences between these two optimal allocations, which reflect the severity of our model’s externality, in Section 4. Finally, the optimal allocations yield a financial network of interest, whose edges represent direct exposures between banks.

3.1 Financial Frictions

It is essential to note that the fundamental financial friction in this model is that markets are incomplete. Bank i is unable to fully hedge the risk associated with another bank j ’s distress process N_t^j . A bank may only partially insure against this risk by holding capital in the risk-free bond. In contrast, however, bank i can control the risk of their *own* jump process N_t^i by choosing to hold cash holdings. However, bank i ’s motivation for doing so is entirely self-preservational. They are not punished for the negative externality inflicted on other firms in the system because: i) there is no contract between banks allowing them to transfer the risk of a jump, and ii) bank i is itself never responsible for losses to external investors. Given that such a friction exists, it is not a novel result that the equilibrium obtained by individual agents operating selfishly is distinct from the ‘welfare-maximizing’ equilibrium obtained by a social planner.

3.2 Decentralized Network

Let us define the value function of bank i to be the supremum over all admissible controls of their expected utility at the terminal time:

$$V_i(t, x) = \sup_{(c^i, w^i) \in \mathcal{A}_{t,T}^i} \mathbb{E} [U_i(X_T^i) | X_t^i = x]. \quad (3.1)$$

Recall that $\mathcal{A}_{t,T}^i$ denotes the set of admissible controls for bank i – defined in Section 2.2. Note also that each bank simultaneously solves their own optimization problem, and therefore the value function of bank i in (3.1) may depend on other banks’ capital allocations. In this sense, the decentralized setting reflects a stable market equilibrium.

Our first result derives the non-local dynamic programming (often referred to as the Hamilton-Jacobi-Bellman or HJB) equation for the value function under regularity.

Proposition 3.1. *If there exist optimal controls and the value function in (3.1) is $\mathcal{C}^{1,1}([0, T], \mathbb{R}_+)$,*

then it solves the following non-local partial differential equation (PDE):

$$0 = \partial_t V_i + \sup_{c_i, w_i} \left\{ \left[(1 - c_i) r + \sum_{j \neq i} w_{ij} \mu_j + \frac{\eta_i \mu_i}{\phi_i} \right] x \partial_x V_i + \theta_i \bar{F}_i(c_i) \left[V_i(t, x(1 - \eta_i)) - V_i \right] \right. \\ \left. + \sum_{j \neq i} \theta_j \bar{F}_j(c_j) \left[V_i(t, x(1 - \phi_j w_{ij})) - V_i \right] \right\}, \quad (3.2)$$

with terminal condition $V_i(T, x) = U_i(x)$. Where unspecified, the value function and its derivatives are evaluated at (t, x) .

The proof is contained in Appendix A.1, and follows from applying Itô's formula to the value function between t and an appropriately defined sequence of stopping times. The assumption that optimal controls exist is verified by a subsequent result in Corollary 3.3.

Fortunately, it is possible to find a separable solution to (3.2), and explicit solutions for the optimal allocations. It is convenient to introduce the following notation:

$$\Gamma(\delta; \gamma) = \begin{cases} \frac{1 - (1 - \delta)^{1 - \gamma}}{1 - \gamma} & \gamma > 0, \gamma \neq 1 \\ -\log(1 - \delta) & \gamma = 1, \end{cases} \quad (3.3)$$

for any $\delta \in [0, 1)$. Within this range, we note that $\Gamma \geq 0$. There is a natural interpretation of this object; for a utility function of the form in (2.3), $\Gamma(\delta; \gamma)$ is proportional to the loss in utility caused by losing a fraction δ of one's wealth. More precisely, $\Gamma(\delta; \gamma_i) = x^{\gamma_i - 1} [U_i(x) - U_i(x(1 - \delta))]$ for any $x > 0$.

We can now state our second main result, which presents a solution to (3.2) and computes the optimal decentralized allocation of capital for bank i .

Proposition 3.2. *The unique optimal cash holdings and project investment amounts for the maximization problem in (3.2) are given by*

$$\hat{c}_i = \begin{cases} f_i^{-1} \left(\frac{r}{\theta_i \Gamma(\eta_i; \gamma_i)} \right) & \text{if } \frac{r}{\theta_i \Gamma(\eta_i; \gamma_i)} \leq f_i(0) \\ 0 & \text{otherwise} \end{cases} \quad \forall i = 1, \dots, n \\ \hat{w}_{ij} = \begin{cases} \frac{1}{\phi_j} \left(1 - \left(\frac{\phi_j \theta_j \bar{F}_j(\hat{c}_j)}{\mu_j} \right)^{1/\gamma_i} \right) & \text{if } \frac{\phi_j \theta_j \bar{F}_j(\hat{c}_j)}{\mu_j} < 1 \\ 0 & \text{otherwise.} \end{cases} \quad \forall j \neq i. \quad (3.4)$$

Furthermore, with the notation

$$J_i^* = \frac{\eta_i \mu_i}{\phi_i} + (1 - \hat{c}_i) r - \theta_i \bar{F}_i(\hat{c}_i) \Gamma(\eta_i; \gamma_i) + \sum_{j \neq i} \hat{w}_{ij} \mu_j - \theta_j \bar{F}_j(\hat{c}_j) \Gamma(\phi_j \hat{w}_{ij}; \gamma_i),$$

the following are explicit solutions to (3.2):

- (i) if $\gamma_i = 1$ and $U_i(x) = \log x$, we have $V_i(t, x) = g_i(t) + \log x$, where $g_i(t) = (T - t) J_i^*$
- (ii) for $\gamma_i \neq 1$ and $U_i(x) = \frac{x^{1 - \gamma_i}}{1 - \gamma_i}$, then we have $V_i(t, x) = g_i(t) U_i(x)$, where $g_i = e^{(1 - \gamma_i)(T - t) J_i^*}$.

The proof, which is again given in Appendix A.1, follows from plugging in the proposed solution, simplifying, and then analyzing the necessary and sufficient conditions for optimality of the resulting maximization problem. A key observation in this proof is that the maximization problem in (3.2) is additively separable between each of the controls c_i, w_i .

Remark 3.1. *The optimal interbank investment \hat{w}_{ij} depends explicitly on \hat{c}_j through the function $\bar{F}_j(\hat{c}_j)$. Moreover, for any choice of bank j 's cash holdings, there exists a corresponding optimal value of w_{ij} . In a game-theoretic sense, this would be bank i 's best response to j 's decision. However, bank j 's optimal value of \hat{c}_j depends only on fixed model parameters. This ensures that \hat{c}_j is bank j 's best response to any decisions made by the other banks, and is therefore a dominant strategy. Hence, the 'game' is trivialized – one can compute every other banks' optimal \hat{c}_j , after which the corresponding \hat{w}_{ij} 's can be easily found.*

The final result of this subsection verifies that the solution given in Proposition 3.2 is indeed equal to the value function.

Corollary 3.3. *The value function in (3.1) is given by*

$$V_i(t, x) = \begin{cases} g_i(t) + \log x & \text{if } \gamma_i = 1 \\ g_i(t) \frac{x^{1-\gamma_i}}{1-\gamma_i} & \text{otherwise,} \end{cases}$$

where $g_i(t)$ and the optimal controls are given in Prop. 3.2.

The proof in Appendix A.1 uses a verification argument. We show that any solution to (3.2) that is once continuously differentiable in both time and space is equal to the value function. Since the proposed solutions in Proposition 3.2 satisfy this regularity condition, we conclude the desired claim. Finally, we note that this result verifies the assumption made in Proposition 3.1 regarding the existence of optimal controls.

Analysis of Decentralized Equilibrium

In analyzing our results, we first note that each unit of additional cash provides some quantifiable marginal benefit by lowering the risk of a liquidity shortage. The quantity $\Gamma(\eta_i, \gamma_i)$, defined in (3.3), is closely related to the loss in utility when a bank with risk aversion γ_i loses a fraction η_i of their wealth. However, the instantaneous hazard rate of an event in which i suffers this loss in utility is controlled by i 's cash holdings through the intensity $\theta_i \bar{F}_i(c_i)$ of the thinned Poisson process. From the proof of Proposition 3.2, the optimal choice of \hat{c}_i will solve the following:

$$\max_{c_i \geq 0} \{ -rc_i - \theta_i \bar{F}_i(c_i) \Gamma(\eta_i; \gamma_i) \},$$

which indicates that the resulting \hat{c}_i achieves the optimal tradeoff between the cost of maintaining liquidity reserves versus the instantaneous expected losses in utility due to a shortage. In particular, the optimal \hat{c}_i from the first-order condition of $r = -\theta_i f_i(\hat{c}_i) \Gamma(\eta_i; \gamma_i)$ ensures that the marginal cost of holding a single additional unit of liquidity (r) equals the marginal reduction in expected losses provided by that unit of liquidity ($\theta_i f_i(\hat{c}_i) \Gamma(\eta_i; \gamma_i)$). In the extreme case where r is large, it may be too costly (relative to the potential losses) for a bank to hold any amount of cash holdings, leading to $\hat{c}_i = 0$.

With explicit solutions for the optimal allocations given in Proposition 3.2, it is possible to analyze their dependence on the exogenous parameters of the system. Due to the monotonicity assumption on $f_i(\cdot)$ in Assumption 1, we can see that the optimal cash holdings \hat{c}_i are decreasing in the liquidity risk premium r : the greater the opportunity cost of holding liquidity, the less cash holdings is held by banks. We also see that $\Gamma(\eta_i; \gamma_i)$ is increasing in η_i . It follows that bank i 's optimal cash holdings, \hat{c}_i increase with the fraction of wealth they stand to lose if a liquidity shortage occurs. In Section 4.2 we will discuss how a regulator can achieve a more socially-optimal allocation by stipulating new values for η_i , each bank's degree of self-investment.

In addition, the optimal project investment by bank i in project j , \hat{w}_{ij} , depends on i only through their risk aversion parameter γ_i . Hence, if $\gamma_i = \gamma_k$ then $\hat{w}_{ij} = \hat{w}_{kj}$. Although the fractional amount of these investments are equal, the nominal amounts may differ because banks i and k may have different levels of wealth. However, the optimal investment amount is decreasing in the investing bank’s risk aversion coefficient γ_i , as we might expect.

The quantity $\frac{\mu_j}{\phi_j \theta_j \bar{F}_j(\hat{c}_j)}$, which appears in (3.4) for \hat{w}_{ij} , is similar to the well-known Sharpe ratio. However, there is one main difference: the variance of bank j ’s project returns can be controlled by j itself. Nonetheless, notice that the optimal investment \hat{w}_{ij} grows with this ‘Sharpe-like’ ratio. In particular, note that if $\frac{\mu_j}{\phi_j \theta_j \bar{F}_j(\hat{c}_j)} < 1$, bank i would prefer to short project j . Since this is not permitted in our model, bank i resorts to an investment of zero. As a direct result, notice that network’s sparsity structure is dictated by this quantity – a project j has external investors if and only if $\frac{\mu_j}{\phi_j \theta_j \bar{F}_j(\hat{c}_j)} > 1$. This implies a ‘core-periphery’ network structure, such that only a subset of banks’ projects are invested in – an example of such a financial network can be seen in Figure 3.1.




Figure 3.1: Sample financial network generated by the decentralized optimum. Edges point from investing banks to projects, and their width indicates the nominal size of the exposures. Node size indicates to total capitalization.

Table 3.1: Parameters for the financial network in Figure 3.1. The risk-free rate is $r = 5\%$. Relevant code can be found here.

Bank	μ_i (%)	ϕ_i	η_i	θ_i	γ_i	$\bar{F}_i(x)$
1	0.9	0.2	0.5	0.04	0.5	$e^{-0.5x}$
2	1	0.3	0.6	0.08	1.7	$e^{-0.6x}$
3	1.5	0.9	0.7	0.12	1	$e^{-0.7x}$
4	1.3	0.6	0.4	0.05	0.3	e^{-2x}
5	1.3	0.82	0.9	0.02	0.87	$e^{-2.4x}$

The resulting ‘core-periphery’ structure reflects the fact that in equilibrium, only a subset of

banks' projects have expected *excess* return greater than zero. The expected return associated with the jump term in Eq. (2.1) is equal to $-\phi_i\theta_i\bar{F}_i(\hat{c}_i)$, whereas the drift is $\mu_i + r$. Hence, the condition $\frac{\mu_i}{\phi_i\theta_i\bar{F}_i(\hat{c}_i)} > 1$ is equivalent to $\mu_i + r - \phi_i\theta_i\bar{F}_i(\hat{c}_i) > r$, which implies that only banks whose projects have expected return greater than r will attract external investors. This is an intuitive consequence of the model, as otherwise investors would not be rewarded for their exposure to the riskiness of these projects. Although the core-periphery structure is determined only by the expected return of a bank's project, we note that both the volatility and return of a project are impacted through a bank's cash holdings. Namely, the optimal w_{ji} decreases in $\phi_i\theta_i\bar{F}_i(\hat{c}_i)$ due to the effects of *both* lower expected return and increased volatility.

A key feature driving the core-periphery structure is the fact that investing banks are identical up to their risk aversion level, which affects only their decision of how much capital is allocated to a project. However, whether or not bank i chooses to invest in another project j depends *only* on the expected profitability of project j . This reflects a system in which banks operate as 'price-takers' in other projects, without the ability to alter the qualities of their investment opportunities.

3.3 Centralized Network

Consider now the perspective of a single social planner of the financial system. In contrast with Section 3.2, we will see that the planner has two different incentives for maintaining bank i 's cash holdings. The first is identical – bank i stands to lose wealth when they suffer a liquidity shortage. The second incentive is concerned with the entire financial system – other banks face losses to their investments on the very same event. Therefore, we expect the planner to have stronger incentives to hold cash, and elect for a greater supply of liquidity throughout the system.

We assume that the planner seeks to maximize social welfare in the system – defined as the sum of all banks' utilities. The planner's value function is therefore given by the following:

$$V(t, x_1, \dots, x_n) = \sup_{(c_i, w_i) \in \mathcal{A}_{t,T}} \mathbb{E} \left[\sum_{i=1}^n U_i(X_T^i) \middle| (X_t^1, \dots, X_t^n) = (x_1, \dots, x_n) \right], \quad (3.5)$$

where $\mathcal{A}_{t,T} = \prod_i \mathcal{A}_{t,T}^i$ is the Cartesian product of each bank's admissible controls.

Remark 3.2. *It is important to note that there are many possible social welfare functions for the planner to consider. In this section, we will see that using the sum of utilities allows for separable solutions to the value function when all banks have logarithmic utility, i.e. $\gamma_i = 1$ for all i . We note that if the planner maximized the product of utilities, then we can also find an explicit solution and optimal controls in the case where $\gamma_i \in (0, 1)$ for all i , but we omit these calculations for conciseness.*

We can relate the planner's value function to those of individual banks in (3.1). The optimal decentralized allocation from Section 3.2 is always feasible for the planner, and therefore the planner's value function is bounded from below by the sum of each bank's value function as follows:

$$V(t, x_1, \dots, x_n) \geq \sum_{i=1}^n V_i(t, x_i). \quad (3.6)$$

This inequality reflects a potential inefficiency of the decentralized setting; the planner's allocation is the first-best (i.e. welfare-maximizing) outcome for the system, and decentralized banks may not achieve this outcome. However, it is not yet clear that the inequality in Eq (3.6) is strict. In Section 4, we will show that in most cases the decentralized optimum is not welfare-maximizing, and establish technical conditions under which we can explicitly compute the gap in welfare.

In what follows, we analyze the planner's optimal allocation by deriving the dynamic programming equation and analyzing the resulting optimization problem. As in the previous section, we first derive the non-local PDE solved by the planner's value function.

Proposition 3.4. *If there exist optimal controls, and the value function in (3.5) is $\mathcal{C}^{1,1,\dots,1}([0, T], \mathbb{R}_+, \dots, \mathbb{R}_+)$, then it solves*

$$0 = \partial_t V + \sup_{c, w} \left\{ \sum_{i=1}^n \left(\left[(1 - c_i) r + \sum_{j \neq i} w_{ij} \mu_j + \frac{\eta_i \mu_i}{\phi_i} \right] x_i \partial_{x_i} V + \theta_i \bar{F}_i(c_i) \left[V(t, x_1(1 - \phi_i w_{1i}), \dots, x_i(1 - \eta_i), \dots, x_n(1 - \phi_i w_{ni})) - V \right] \right) \right\}, \quad (3.7)$$

with terminal condition $V(T, x_1, \dots, x_n) = \sum_{i=1}^n U_i(x_i)$. Where unspecified, the value function and its derivatives are evaluated at (t, x_1, \dots, x_n) .

The proof of this result can again be found in Appendix A.2. Proposition 3.4 yields an $n + 1$ dimensional non-local PDE for the planner's value function. There is one key difference between Equations (3.7) and (3.2) – when a bank suffers a liquidity shortage, the planner's value function is affected by losses occurring throughout the *entire* financial system. This is not true in the decentralized setting; an individual bank's value function only depends on their own losses caused by such an event.

With specific choices of utility functions, it is possible to find a separable solution to (3.7), and prove existence of an optimal allocation. However, in order to establish uniqueness, we will need the following technical assumption.

Assumption 2. *Let each shock density $f_i(\cdot)$ satisfy*

$$\frac{f_i(c)}{\bar{F}_i(c)} + 3 \frac{f_i'(c)}{f_i(c)} - \frac{f_i''(c)}{f_i'(c)} < 0, \quad \forall c \geq 0. \quad (3.8)$$

and, with the notation $\tilde{c}_i = F_i^{-1} \left(\left[1 - \frac{\mu_i}{\phi_i \theta_i} \right]_+ \right)$, assume that the following holds for all i

$$\Gamma(\eta_i; 1) > \begin{cases} \min \left\{ (n-1) \left[\log \left(\frac{\phi_i \theta_i}{\mu_i} \right) - \frac{f_i(0)^2}{f_i'(0)} \right], \frac{r}{\theta_i f_i(0)} + (n-1) \log \left(\frac{\phi_i \theta_i}{\mu_i} \right) \right\} & \text{if } \tilde{c}_i = 0 \\ \min \left\{ -(n-1) \frac{\phi_i \theta_i f_i(\tilde{c}_i)^2}{\mu_i f_i'(\tilde{c}_i)}, \frac{r}{\theta_i f_i(\tilde{c}_i)} \right\} & \text{otherwise,} \end{cases} \quad (3.9)$$

where recall from Eq. (3.3) that $\Gamma(\eta_i; 1) = \log(1) - \log(1 - \eta_i) = -\log(1 - \eta_i)$ represents the loss in utility associated with a loss of η_i fraction of wealth.

Assumption 2 provides sufficient – but not necessary – conditions for uniqueness of the planner's optimal allocation. Numerically, we have observed that the optimal solution is often unique, but the optimization problem in (3.7) is (generally) not convex, and therefore proving uniqueness is non-trivial. We do, however, note that the inequality (3.8) is always satisfied by exponential and power distributions. At a glance, Eq. (3.9) requires that each bank's losses upon a liquidity shortage not be too small. If bank i loses only a small amount of their wealth when a liquidity shortage occurs, then it is possible for there to be two equilibria for the planner: i) project i is profitable and bank i 's cash holdings are high; and ii) project i is not profitable and bank i 's cash holdings are low. Since the same rationale holds for all n projects, we could have up to 2^n distinct equilibria.

In this paper, we will focus on settings where unique optimal controls can be provably obtained, and hence require Assumption 2 to hold.

Analogous to Section 3.2, we show there exists a separable solution to the PDE (3.7). Additionally, we show that the optimal solution will solve a system of algebraic equations.

Proposition 3.5. *Let each bank have a logarithmic utility function (i.e. $\gamma_i = 1 \forall i$). Then, there exist optimal cash holdings and project investment amounts for the planner, which solve the following system of equations:*

$$\begin{aligned} c_i^* &= \begin{cases} f_i^{-1} \left(\frac{r}{\theta_i [\Gamma(\eta_i; 1) + (n-1)\Gamma(\phi_i w_i^*; 1)]} \right) & \text{if } f_i(0) \leq \frac{r}{\theta_i [\Gamma(\eta_i; 1) + (n-1)\Gamma(\phi_i w_i^*; 1)]} \\ 0 & \text{otherwise,} \end{cases} \\ w_i^* &= \begin{cases} \frac{1}{\phi_i} \left(1 - \frac{\phi_i \theta_i \bar{F}_i(c_i^*)}{\mu_i} \right) & \text{if } \frac{\phi_i \theta_i \bar{F}_i(c_i^*)}{\mu_i} < 1 \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \quad (3.10)$$

We define

$$J_C^* = \sum_{i=1}^n \left(\left[(1 - c_i^*) r + (n-1) w_i^* \mu_i + \frac{\eta_i \mu_i}{\phi_i} \right] - \theta_i \bar{F}_i(c_i^*) \left[\Gamma(\eta_i; 1) + (n-1)\Gamma(\phi_i w_i^*; 1) \right] \right), \quad (3.11)$$

and $g(t) = (T - t) J_C^*$. The solution to (3.7) is given by

$$V(t, x_1, \dots, x_n) = g(t) + \sum_{i=1}^n \log x_i. \quad (3.12)$$

Furthermore, under Assumption 2, the optimal controls (c_i^*, w_i^*) are unique.

The proof is again given in Appendix A.2. We note that a separable solution using logarithmic utility functions is only possible because the planner aims to maximize the sum of banks' expected utilities. See Remark 3.2 for a brief discussion of other settings where a separable solution can be obtained.

In contrast to the decentralized setting, the maximization in (3.7) is not additively separable between each optimization variable. Nonetheless, each of the i subsets $\{c_i, w_{1i}, \dots, w_{ni}\}$, $i = 1, \dots, n$ can be analyzed separately, which greatly simplifies our analysis. However, the coupling between c_i and w_i leads to the need for additional assumptions to establish uniqueness.

The system of equations in (3.10) admits a block coordinate descent approach. Namely, for any fixed c_i , the maximization problem for w_i is strictly concave and admits a unique solution (these can be seen in the proof of Proposition 3.5). Conversely, for given values of w_i , the maximizing of c_i shares these features. As a result, we can iteratively update these variables to solve for the planner's optimum numerically. Upon convergence, we are guaranteed to have found the uniquely optimal allocation.

Since we have shown existence of an optimal allocation, and the proposed solution in (3.12) is continuously differentiable, then we are able to verify that it is indeed equal to the planner's value function.

Corollary 3.6. *The planner's value function in (3.5) is given by (3.12). Furthermore, the optimal project investment and cash holdings solve (3.4).*

As with previous results, the proof is presented in Appendix A.2.

Analysis of Centralized Equilibrium

There is one main difference between the system of equations in (3.10) and the optimal solutions from the decentralized setting in (3.4). Here, the expression for the planner's optimal cash holdings c_i^* contains an additional term of $(n-1)\Gamma(\phi_i w_i^*; 1)$. This term directly captures the externality – when bank i suffers a liquidity shortage and their project incurs losses, the planner is cognizant of losses in utility experienced by all banks. As a result, with more banks the planner maintains larger cash holdings to compensate for greater system-wide losses. In contrast, bank i 's decentralized optimization problem considers only changes to their own wealth, and therefore their optimal \hat{c}_i will be indifferent to the system's size.

Since we will have $w_i^* \geq 0$ in (3.10), the planner has no weaker an incentive to hold cash than an individual bank.² Hence, the planner will hold no less cash than banks in the decentralized setting – this difference will be studied more closely in the following section. Finally, we also notice that given some fixed cash holdings, the optimal project investments w_i^* and \hat{w}_i are computed identically in both settings. It follows that any differences between the optimal project investments in (3.4) and (3.10) must be driven only by differences in optimal cash holdings, which we proceed to analyze in Section 4.

Finally, we note that *only* in the case where w_i^* equal zero for all i will the planner's optimal allocation coincide with that of the decentralized system. In such a case, bank i 's cash holdings are dictated only by their own losses, which implies that $\hat{c}_i = c_i^*$. In other words, as long as there is least one project within the planner's optimum that has positive expected excess return, the market-mediated equilibrium from Section 3.2 is inefficient, and the inequality in Eq. (3.6) is strict. This suggests that there is value to be obtained from regulating the system.

4 Differences in Equilibria

It is natural to compare the two optimal allocations from Sections 3.2 and 3.3. In particular, we may be interested in computing the gap in welfare from the inequality (3.6). Figure 4.1 illustrates a sample path for the wealth of three banks, where in Fig. 4.1a the controls are given by (3.4), and in Fig. 4.1b by (3.10). Qualitatively, there are higher-frequency jumps in 4.1a, but the jumps are of larger size in 4.1b. With the remainder of this section, we study these differences more precisely.

In what follows, we will assume that all banks have logarithmic utility (i.e. $\gamma_i = 1$ for all i). Recall that \hat{c}_i, \hat{w}_{ji} denote the optimal decentralized allocations given in (3.4). Note that for all $j, k \neq i$ we will have $\hat{w}_{ki} = \hat{w}_{ji}$, so we will denote these fractional amounts to be \hat{w}_i (this follows from $\gamma_j = 1$ for all j). Additionally, recall that c_i^*, w_i^* denotes the optimal solution from (3.10). Finally, we use the asymptotic notation $g(n) = \Theta(h(n))$ to denote that there exist positive constants A_1, A_2 such that $A_1 \leq \lim_{n \rightarrow \infty} \frac{g(n)}{h(n)} \leq A_2$. If $A_1 = A_2$, then we will write $g(n) \asymp h(n)$.

Comparing the two optimal allocations, since $w_i^* \geq 0$, it will necessarily be the case that $c_i^* \geq \hat{c}_i$. Our first core result establishes the asymptotic rate at which the planner's optimal cash holdings grow as the system's size increases, when liquidity shocks follow an exponential distribution. More precisely, we show that the planner's cash holdings must grow at logarithmically in the system size n . In contrast, if $w_i^* = 0$, then we would have $c_i^* = \hat{c}_i$, which is of constant order.

²This observation may not be the case if, for example, short-selling were allowed. Qualitatively, the planner may choose to have a single bank i hold zero cash, while others in the system maintain large, short positions in i 's project. In this case, the total utility of the system may actually increase when bank i 's project incurs losses. However, clearly this result may not align with the best outcome for bank i itself.

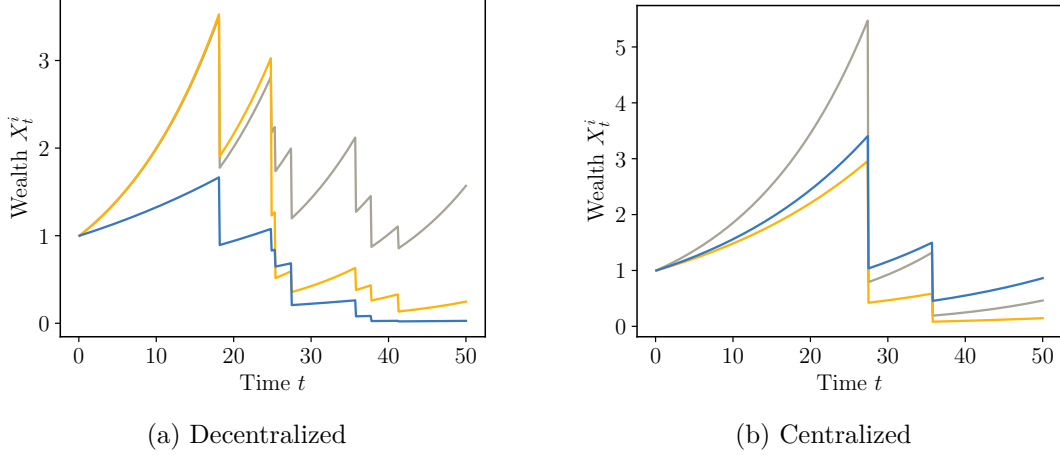


Figure 4.1: An example of wealth dynamics under the both optimal allocations for a system of $n = 3$ banks. The same random seed is used in both simulations, so that the size and arrival times of liquidity shocks are identical. For conciseness, we do not include the parameters, but the code to reproduce these figures can be found here.

Proposition 4.1. *If $F_i(x) = 1 - e^{-\frac{x}{\lambda_i}}$, then*

$$\lambda_i \log \left(\frac{\theta_i(n-1)}{\lambda_i r} \left[\log(n-1) - \log \left(\frac{\Gamma(\eta_i; 1)}{\Gamma(\phi_i \hat{w}_i; 1)} \right) \right] \right) \leq c_i^* \leq \lambda_i \log \left(\frac{\theta_i C_U(n-1)}{\lambda_i r} \log(n) \right),$$

where $\Gamma(\cdot; 1)$ is defined in Eq. (3.3).

In particular, $c_i^* \asymp \lambda_i \log(n)$.

The proof follows from Lemma A.1 in Appendix A, which establishes upper and lower bounds on super- and sub-exponentially distributed liquidity shocks.

This result highlights a substantial discrepancy between the planner's optimum and that of individual banks. Since the losses to an individual bank are not affected by the number of external investors, their cash holdings are dictated only by their own losses and are of constant order. In contrast, the planner observes losses throughout the system, which grow at least linearly in the system size, and hence chooses to hold greater cash amounts to compensate for the greater losses upon a liquidity shortage to any particular bank.

Proposition 4.1 is a useful tool for comparing the two optimal allocations, as all differences are driven by the distinct cash holdings. The following result compares the asymptotic rates for various quantities in both the centralized and decentralized equilibrium.

Corollary 4.2. *Under the conditions of Proposition 4.1, and when $\hat{w}_i > 0$, we have:*

(i) **Likelihood of Liquidity Shortage:**

$$\bar{F}_i(c_i^*) = \Theta \left(\frac{1}{n \log(n)} \right), \quad \bar{F}_i(\hat{c}_i) = \Theta(1).$$

(ii) **Optimal Project Investment:**

$$w_i^* = \frac{1}{\phi_i} - \Theta \left(\frac{1}{n \log(n)} \right), \quad \hat{w}_i = \Theta(1).$$

(iii) *Expected System-wide Loss in Utility (due to bank i):*

$$\bar{F}_i(c_i^*)(n-1)\Gamma(w_i^*;1) = \Theta(1), \quad \bar{F}_i(\hat{c}_i)(n-1)\Gamma(\hat{w}_i;1) = \Theta(n).$$

The proofs require only plugging in the asymptotic rate for c_i^* from Proposition 4.1 into $\bar{F}_i(\cdot)$, the expression for w_i^* , and $\Gamma(w_i;1)$.

Notice that we must have $w_i < \phi_i^{-1}$ to ensure wealth remains positive, yet we can still pin down the rate at which the optimal project investments approach their upper bound. The term $\bar{F}_i(c_i^*)(n-1)\Gamma(w_i^*;1)$ appears in Eq. (3.11), the value function for the planner. It represents the expected loss in utility to all of bank i 's external investors when i suffers a liquidity shortage: $\Gamma(w_i^*;1)$ captures the loss in utility to a single investing bank, of which there are $n-1$ within the system, and $\bar{F}_i(c_i^*)$ is proportional to the likelihood of such an event. It is particularly interesting that this quantity is of constant order, as it shows that the planner perfectly compensates for larger possible losses of utility in the system by reducing the probability of such an event.

The results in Corollary 4.2 allow us to analyze differences in welfare between the two settings in the following section.

4.1 Price of Anarchy

We now turn to the gap between value functions from (3.6). It will be useful to have \mathcal{M}_n denote the set of banks whose projects are invested in by other firms in the planner's optimal allocation, i.e. $\mathcal{M}_n = \{i \in \{1, \dots, n\} : w_i^* > 0\}$. Banks in \mathcal{M}_n form the financial network's 'core'. If for some i we have $w_i^* = 0$, then it must be the case that $c_i^* = \hat{c}_i$ and $\frac{\phi_i \theta_i \bar{F}_i(c_i^*)}{\mu_i} > 1$. For such a bank i , the planner's optimal c_i^* would remain constant at \hat{c}_i , even as n grows.

In this model, we measure this efficiency gap through the price of anarchy, defined as:

$$\text{PA} = \frac{V}{\sum_{i=1}^n V_i}.$$

This quantity reflects the relative gap in social welfare due to greedy decentralized behavior in the system (Papadimitriou 2001). In the following result, we characterize its asymptotic behavior.

Proposition 4.3. *Assume that $\gamma_i = 1$ and $F_i(x) = 1 - e^{-\frac{x}{\lambda_i}}$ for all i . Then, as $n \rightarrow \infty$, we have*

$$\text{PA} = \Theta(1).$$

The proof is found in Appendix A.3, and uses the results from Corollary 4.2.

It is particularly interesting that the price of anarchy does not grow with the system size n , nor the remaining time horizon $(T-t)$. A more precise result can be obtained if banks are sufficiently homogeneous, where we can compute the limiting value.

Corollary 4.4. *In the setting of Proposition 4.3, assume further that all banks in \mathcal{M}_n are identical (i.e. $\mu_j = \mu$, $\phi_j = \phi$, $\theta_j = \theta$, $\eta_j = \eta$, and $\lambda_j = \lambda$ for some given constants μ, ϕ, θ, η and λ). If $|\mathcal{M}_n| \xrightarrow[n \rightarrow \infty]{} \infty$, then*

$$\frac{V_i}{|\mathcal{M}_n|(T-t)} \xrightarrow[n \rightarrow \infty]{} \frac{\mu}{\phi} + \theta \bar{F}(\hat{c}) \left[\log \left(\frac{\phi \theta \bar{F}(\hat{c})}{\mu} \right) - 1 \right], \quad \forall i = 1, \dots, n$$

$$\frac{V}{n|\mathcal{M}_n|(T-t)} \xrightarrow[n \rightarrow \infty]{} \frac{\mu}{\phi}.$$

where \hat{c} is given in (3.4) and $\bar{F}(\hat{c}) = e^{-\frac{\hat{c}}{\lambda}}$. As a result, we have:

$$\text{PA} \xrightarrow{n \rightarrow \infty} \frac{1}{1 + \frac{\phi\theta\bar{F}(\hat{c})}{\mu} \left[\log \left(\frac{\phi\theta\bar{F}(\hat{c})}{\mu} \right) - 1 \right]}. \quad (4.1)$$

Corollary 4.4 verifies that the price of anarchy is of constant order in n , and the proof is found in Appendix A.3. Of particular interest, the rate at which $|\mathcal{M}_n|$ grows in n does not appear in our result. This implies that the limiting value for the price of anarchy is independent from the fraction of the system that operates as its ‘core’. Notice also that $\phi\theta\bar{F}(\hat{c}) < \mu$, and hence the right-hand side in (4.1) is greater than one. Moreover, the limiting value is increasing in $\frac{\phi\theta\bar{F}(\hat{c})}{\mu}$. Therefore, as the profitability of projects in the decentralized setting is reduced, price of anarchy diverges to infinity with n .

Corollary 4.4 is verified numerically. Using the parameters in Table 4.1, we compute the individual and collective value functions. The resulting PA is plotted in Figure 4.2, along with the limiting value in (4.1). We see that this quantity quickly converges to its limit.




Figure 4.2: Simulating the price of anarchy for a system of identical firms as n grows. Relevant parameters are shown in Table 4.1.

Table 4.1: System parameters used in simulations for Figures 4.2 and C.1. Relevant code is available here.

Notation	Value	Description
r	0.01	Risk-free rate
μ	0.045	Excess drift
ϕ	0.4	Losses to External Investors
η	0.5	Losses to Associated Bank
$F(x)$	$1 - e^{-\frac{x}{\lambda}}$	CDF of shock size
λ	1	Parameter of $F(\cdot)$
θ	0.1	Shock arrival rate
γ	1	Relative risk aversion coefficient

4.2 Replicating the Centralized Allocation

Finally, we may be interested in studying how banks in the decentralized setting can be incentivized or forced to replicate the planner’s optimal allocation. We could approach this regulatory question in one of two ways – either directly enforcing a liquidity requirement on individual banks (corresponding to the planner’s optimal c_i^*), or allowing regulators to control the previously exogenous parameter η_i , which dictates how much of bank i ’s wealth must be invested in their own project. In this section, we choose to pursue the latter approach.

The rationale for this is twofold: first, η_i plays the fundamental role in decentralized banks’ choice of optimal cash holdings. From the analysis near the end of Section 3.2, we observed that by increasing η_i , banks will increase their cash holdings and therefore reduce their projects’ riskiness, which can lead them to eventually match the planner’s optimal liquidity reserves. Second, we can imagine that externally investing banks are permitted to write a contract stipulating the associated bank’s degree of co-investment. This kind of contracting is not a focus of our paper, but is instead analyzed in more detail with Principal-Agent problems such as Hernández Santibáñez et al. (2020). Nonetheless, the co-investment contract can be designed to ensure that individual banks hold sufficient cash.

There are several other strategies used by regulators in practice to reduce inefficiencies in the financial system, such as bail-in or bail-out policies, or government intervention (Bernard et al. 2022, Kanik 2019). While these strategies are fruitful directions for future work, we will focus on how co-investment requirements can be used to reduce the riskiness of core banks in the system, since these parameters are already integrated into the model and directly imply equilibrium levels of cash holdings.

Let η_i^C (respectively η_i^D) denote the fraction of bank i ’s wealth lost upon liquidity shortage in the centralized (resp. decentralized) setting. We would like to choose η_i^D so that decentralized banks replicate the planner’s optimum with given values η_i^C . More precisely, we seek to find η_i^D solving $c_i^*(\eta_i^C) = \hat{c}_i(\eta_i^D)$ for all i , where we write the optimal controls in a way that highlights their dependence on the underlying values of η . Even though the optimal allocations are identical, however, we note that the decentralized optimum is still inefficient (with respect to the optimal centralized allocation corresponding to η_i^D). Using equations (3.4) and (3.10), we find that:

$$\eta_i^D = 1 - (1 - \eta_i^C) (1 - \phi_i w_{.i}^*(\eta_i^C))^{n-1}.$$

First, notice that whenever $w_{.i}^*(\eta_i^C) > 0$, the resulting value of η_i^D will grow with n towards its upper bound of 1. This is intuitive – banks whose projects are highly invested in require the strongest incentive to reduce their project’s riskiness. It is precisely these banks that comprise the financial network’s ‘core’, and will be subjected to the strictest regulations. In contrast, banks in the periphery (with $w_{.i}(\eta_i^C) = 0$) are no more strictly regulated ($\eta_i^D = \eta_i^C$), since their liquidity shortages cause only localized losses. Second, we see that for banks to replicate the planner’s optimum, it is necessary for bank i ’s degree of co-investment to depend on their liabilities throughout the system. It is therefore necessary to know the complete structure of the financial network to determine the value of η_i^D , which may not be known to individual investing banks. Finally, an interesting case occurs when we choose $\eta_i^C = 0$. In this case, the value of η_i^D is only non-zero if $w_{.i}^*(0) > 0$. Namely, banks in the periphery would not be required to hold any stake in their own projects.

This discrepancy suggests that while stricter regulations on core banks can improve the efficiency of the financial system, there is little if any value to be gained by increasing restrictions on banks in the periphery. When increasing the self-investment requirement of some bank in the periphery, one of two will occur: i) this increase causes i to have large enough cash holdings for other banks

in the system to invest in project i , or ii) the increase in η_i is not sufficiently large for i 's project to be profitable for other banks in the system. In case i), bank i moves into the ‘core’ of the financial system, as their project is now capitalized by other firms. In this case, it is not possible to conclusively state whether or not this increase in η_i yielded higher or lower social welfare. However, in case ii), the small increase to η_i serves only to push bank i 's capital towards less efficient investments – they are obligated to hold additional capital in a project whose expected return is less than r , and they choose to hold greater cash amounts to compensate for this heightened exposure – which also provides lower returns than r . In this case, both effects serve only to reduce i 's wealth and hence the total welfare in the financial system.

5 Conclusion

In this paper, we present a model in which banks in a financial system control both their own levels of cash holdings, and their investment in each others’ risky projects. We compute the unique optimal allocations of capital for two distinct organizations of the system, and study their differences qualitatively and quantitatively. First, we analyze the setting where each bank acts with pure self-interest. We compute explicitly the optimal allocation, and find that the size of project investments are closely related to a Sharpe-like ratio – which is controlled by the bank associated with a particular project. In particular, the equilibrium financial network exhibits a ‘core-periphery’ structure, in which only a subset of banks serve as valid investment targets. Second, we formulate the optimization problem of a social planner, who seeks to maximize the total welfare in the system. Under a few technical assumptions, we are able to prove the existence of a unique optimal allocation. A substantial technical contribution of this paper is to provide explicit solution of a multi-player game of intensity control in a CRRA-utility maximization setting. We also able to solve explicitly the associated social planner’s problem, and characterize the asymptotics of the system’s inefficiencies. Finally, we can also extend the model in various directions described in the Appendices, which can be solved up to numerical solution of systems of algebraic equations.

In particular, we find that the planner’s optimum exhibits low-frequency and high-severity events of distress, which aligns with the ‘robust-yet-fragile’ feature observed by Gai and Kapadia (2010). The difference in these two optimal allocations is driven by a negative externality, where individual banks are excessively risky given the potential losses that they may induce.

In the case where shocks are exponentially distributed, we can precisely compute how the externality’s severity depends on the system’s size. We see that the planner compensates for an increased number of investors by reducing the risk of a bank’s project. The planner perfectly balances the two effects, so that the expected losses in utility remain of constant order – regardless of the system size. We are also able to see that the absolute loss in welfare due to decentralized behavior grows with the size of both the financial system and its core. However, and perhaps counterintuitively, the price of anarchy, is of constant order. This result may be a consequence of the model, wherein firms’ decisions are fractions of their wealth. Finally, we show that it is possible, through regulation or contracting between banks, to replicate the planner’s optimal interbank allocation.

Banks in the system’s core must be subjected to the strictest requirements, and will therefore have the strongest incentive to reduce their project’s riskiness. This highlights both i) the danger of government bailouts, which can cause perverse incentives for individual banks, and ii) the heterogeneous impact of regulation on social welfare.

In the wake of the regional banking crisis in 2023, the widespread financial distress was addressed with forms of bailouts and takeovers of, for instance, Silicon Valley Bank and First Republic Bank.

Subsequently, regulators are discussing revisions to capitalization and liquidity requirements of banks contained in Basel III (Basel Committee on Banking Supervision 2017), which led to a debate and substantial lobbying from large, central banks such as JP Morgan. This resulted in US regulators lowering their initial proposed increases to capital and liquidity requirements. Our work strongly supports the idea of heterogeneous regulation based on a bank’s position within the network.

We believe there are several interesting continuations of this work. First, a notable limitation of this model is that it does not contain a mechanism of contagion. For instance, Aït-Sahalia and Hurd (2015) consider a portfolio optimization problem where assets’ jump components are self- and mutually exciting. An immediate extension of our work may be to incorporate jump processes with these features directly into the model. It may also be possible to show that self- and mutually exciting jumps can endogenously emerge, e.g. if an investing bank suffers losses to their cash holdings when the bank associated with this project suffers. Second, we assume that the liquidity risk premium r is fixed and exogenous, so that banks may borrow as much as they like at this rate. In practice, this rate would be determined endogenously by the relative supply and demand of liquidity, and substantially affected by market conditions. Additionally, financial crises are heavily destabilizing, and it is natural to assume that it is challenging (or impossible) to quickly rebalance a portfolio in the wake of such an event. Therefore, it is practical to prevent banks from instantaneously re-weighting their portfolios. This feature may lead to further inefficiencies caused by banks’ inability to establish an optimal allocation of wealth shortly after a shock occurs. Furthermore, our model differs from the literature on strategic network formation in that creating an ‘investment linkage’ to another bank is costless. It is natural to incorporate these costs into banks’ optimization problems, for example, as there is a cost to performing due diligence on a potential investment opportunity to assess its profitability.

Acknowledgments

We are grateful to Emma Hubert and Roberto Rigobon for helpful comments and discussion.

References

- Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2015a.
- Daron Acemoglu, Asuman E Ozdaglar, and Alireza Tahbaz-Salehi. Systemic Risk in Endogenous Financial Networks. *Columbia Business School Research Paper*, (15-17), 2015b.
- Daron Acemoglu, Asuman Ozdaglar, James Siderius, and Alireza Tahbaz-Salehi. Systemic credit freezes in financial lending networks. *Mathematics and Financial Economics*, 15:185–232, 2021.
- Yacine Aït-Sahalia and T R Hurd. Portfolio Choice in Markets with Contagion. *Journal of Financial Econometrics*, 14(1):1–28, 2015.
- Franklin Allen and Douglas Gale. Financial Contagion. *Journal of Political Economy*, 108(1):1–33, 2000.
- Hamed Amini and Andreea Minca. Inhomogeneous Financial Networks and Contagious Links. *Operations Research*, 64(5):1109–1120, 2016.
- Ana Babus. The Formation of Financial Networks. *The RAND Journal of Economics*, 47(2):239–272, 2016.
- Venkatesh Bala and Sanjeev Goyal. A noncooperative model of network formation. *Econometrica*, 68(5): 1181–1229, 2000.
- Basel Committee on Banking Supervision. Basel iii: international regulatory framework for banks, 2017. URL <https://www.bis.org/bcbs/basel3.htm>.

- Stefano Battiston, Domenico Delli Gatti, Mauro Gallegati, Bruce Greenwald, and Joseph E. Stiglitz. Liaisons Dangereuses: Increasing Connectivity, Risk Sharing, and Systemic Risk. *Journal of Economic Dynamics and Control*, 36(8):1121–1141, 2012.
- Sarah Bensalem, Nicolás Hernández Santibáñez, and Nabil Kazi-Tani. A continuous-time model of self-protection. *hal-02974961v2*, 2020.
- Benjamin Bernard, Agostino Capponi, and Joseph E Stiglitz. Bail-ins and bailouts: Incentives, connectivity, and systemic stability. *Journal of Political Economy*, 130(7):1805–1859, 2022.
- Bruno Biais, Thomas Mariotti, Jean-Charles Rochet, and Stéphane Villeneuve. Large Risks, Limited Liability, and Dynamic Moral Hazard. *Econometrica*, 78(1):73–118, 2010.
- Marcel Bluhm, Ester Faia, and Jan Pieter Krahen. Endogenous Banks’ Networks, Cascades and Systemic Risk. *SAFE Working Paper*, 2014.
- Stephen Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge University Press, 2004.
- Fabio Caccioli, Munik Shrestha, Cristopher Moore, and J. Doyne Farmer. Stability Analysis of Financial Contagion due to Overlapping Portfolios. *Journal of Banking and Finance*, 46(1):233–245, 2014.
- Agostino Capponi and Christoph Frei. Dynamic Contracting: Accidents Lead to Nonlinear Contracts. *SIAM Journal on Financial Mathematics*, 6(1):959–983, 2015.
- Rama Cont and Peter Tankov. *Financial Modelling with Jump Processes*. Chapman and Hall/CRC, 2003.
- Nils Detering, Thilo Meyer-Brandis, Konstantinos Panagiotou, and Daniel Ritter. Managing Default Contagion in Inhomogeneous Financial Networks. *SIAM Journal on Financial Mathematics*, 10(2):578–614, 2019.
- Nils Detering, Thilo Meyer-Brandis, Konstantinos Panagiotou, and Daniel Ritter. Financial Contagion in a Stochastic Block Model. *International Journal of Theoretical and Applied Finance*, 23(08), 2020.
- Nils Detering, Thilo Meyer-Brandis, Konstantinos Panagiotou, and Daniel Ritter. An Integrated Model for Fire Sales and Default Contagion. *Mathematics and Financial Economics*, 15(1):59–101, 2021.
- Matthew Elliott, Benjamin Golub, and Matthew O Jackson. Financial Networks and Contagion. *American Economic Review*, 104(10):3115–53, 2014.
- Selman Erol and Guillermo Ordonez. Network reactions to banking regulations. *Journal of Monetary Economics*, 89:51–67, 2017.
- Maryam Farboodi. Intermediation and Voluntary Exposure to Counterparty Risk. *NBER Working Paper*, (w29467), 2021.
- Prasanna Gai and Sujit Kapadia. Contagion in Financial Networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 466(2120):2401–2423, 2010.
- Prasanna Gai, Andrew Haldane, and Sujit Kapadia. Complexity, Concentration and Contagion. *Journal of Monetary Economics*, 58(5):453–470, 2011.
- Nicolás Hernández Santibáñez, Dylan Possamaï, and Chao Zhou. Bank Monitoring Incentives under Moral Hazard and Adverse Selection. *Journal of Optimization Theory and Applications*, 184(3):988–1035, 2020.
- Matthew O Jackson and Asher Wolinsky. A Strategic Model of Social and Economic Networks. *Journal of Economic Theory*, 71:44–74, 1996.
- Zafer Kanik. From lombard street to wall street: systemic risk, rescues, and stability in financial networks. *Work. Pap., Mass. Inst. Technol., Cambridge*, 2019.
- PA W Lewis and Gerald S Shedler. Simulation of Nonhomogeneous Poisson Processes by Thinning. *Naval Research Logistics Quarterly*, 26(3):403–413, 1979.
- Zongxi Li, A. Max Reppen, and R. Sircar. A mean field games model for cryptocurrency mining. *Management Science*, 70:2188–2208, 2024.
- Robert C Merton. Lifetime Portfolio Selection Under Uncertainty: The Continuous-time Case. *The Review of Economics and Statistics*, pages 247–257, 1969.
- Robert C Merton. Optimum Consumption and Portfolio Rules in a Continuous-time Model. *Journal of Economic Theory*, 3:373–413, 1971.

- Henri Pagès and Dylan Possamai. A Mathematical Treatment of Bank Monitoring Incentives. *Finance and Stochastics*, 1(18):39–73, 2014.
- Christos Papadimitriou. Algorithms, Games, and the Internet. In *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing*, pages 749–753, 2001.
- Carlos A. Ramírez. Regulating financial networks: A flying blind problem. *Under revision*, March 2022.
- Jean-Charles Rochet and Jean Tirole. Interbank Lending and Systemic Risk. *Journal of Money, Credit and Banking*, 28(4):733–762, 1996.
- L. Christopher G. Rogers. *Optimal Investment*. Springer, 2013.
- Adam Zawadowski. Entangled Financial Systems. *The Review of Financial Studies*, 26(5):1291–1323, 2013.

A Proofs

A.1 Decentralized Network

Proof of Proposition 3.1. First, we use the dynamic programming principle to consider only the optimal control over the time interval $[t, \tau]$, for a stopping time $\tau < T$ to be defined later. We can write the value function recursively as

$$V_i(t, x) = \sup_{(c^i, w^i) \in \mathcal{A}_{t,T}^i} \mathbb{E} \left[V_i(\tau, X_\tau^i) \mid X_t^i = x \right], \quad (\text{A.1})$$

which holds for all $t < T$ and $\tau \leq T$.

Next, we for each bank k we fix some admissible control $(c^k, w^k) \in \mathcal{A}_{t,T}^k$. By assumption, V_i is once differentiable in both time and space, and using Itô's formula (see for instance Cont and Tankov (2003)) we can write:

$$\begin{aligned} V_i(\tau, X_\tau^i) - V_i(t, X_t^i) &= \int_t^\tau [\partial_t V_i(s, X_s^i) + \partial_x V_i(s, X_s^i) b_i(c_s^i, w_s^i) X_s^i] ds \\ &\quad + \sum_{j=1}^n \int_t^\tau [V_i(s, X_s^i) - V_i(s, X_{s-}^i)] dN_s^j. \end{aligned} \quad (\text{A.2})$$

where $b_i(c_s^i, w_s^i)$ is the coefficient on the dt term in (2.2).

Recall that the jump process N_t^j has instantaneous intensity $\theta_j \bar{F}_j(c_t^j)$. Therefore, the compensated process $M_t^j = N_t^j - \int_0^t \theta_j \bar{F}_j(c_s^j) ds$ is a martingale. Rewriting the integrals in (A.2) in terms of dM_t^j and taking expectation conditioned on $X_t^i = x$ (denoted $\mathbb{E}_{t,x}$) of both sides yields:

$$\begin{aligned} \mathbb{E}_{t,x} [V_i(\tau, X_\tau^i)] - V_i(t, X_t^i) &= \mathbb{E}_{t,x} \left[\int_t^\tau \mathcal{L}^{c_s^i, w_s^i} V_i(s, X_{s-}^i) ds \right] \\ &\quad + \mathbb{E}_{t,x} \left[\int_t^\tau [V_i(s, X_{s-}^i - \eta_i X_{s-}^i) - V_i(s, X_{s-}^i)] dM_s^i \right] \\ &\quad + \sum_{j \neq i} \mathbb{E}_{t,x} \left[\int_t^\tau [V_i(s, X_{s-}^i - \phi_j w_s^{ij} X_{s-}^i) - V_i(s, X_{s-}^i)] dM_s^j \right], \end{aligned} \quad (\text{A.3})$$

where the generator \mathcal{L}^{c_i, w_i} is defined to be

$$\begin{aligned} \mathcal{L}^{c_i, w_i} \psi(t, x) &= \partial_t \psi(t, x) + \left((1 - c_i) r + \sum_{j \neq i} w_{ij} \mu_j + \frac{\eta_i \mu_i}{\phi_i} \right) x \partial_x \psi \\ &\quad + \theta_i (1 - F_i(c_i)) \left[\psi(t, x(1 - \eta_i)) - \psi(t, x) \right] \\ &\quad + \sum_{j \neq i} \theta_j (1 - F_j(c_j)) \left[\psi(t, x(1 - \phi_j w_{ij})) - \psi(t, x) \right], \end{aligned} \quad (\text{A.4})$$

for any $\psi \in \mathcal{C}^{1,1}([0, T], \mathbb{R}_+)$.

Next, we need to show that the expectation of the stochastic integrals with respect to dM_s^k are equal to zero. To do so, it is sufficient to have the integrand bounded for $s \in [t, \tau]$. Define the stopping time τ to be:

$$\tau = (t + \delta) \wedge \inf \left\{ s \in [t, T], X_s^i \leq \epsilon \text{ or } X_s^i \geq \frac{1}{\epsilon} \right\}, \quad (\text{A.5})$$

for some small $\delta > 0$ and $\epsilon > 0$. Then, since X_s^i is bounded away from zero in $[t, \tau]$, the size in the jump of the value function is bounded. Therefore the stochastic integrals in (A.3) have zero expectation. We obtain:

$$\mathbb{E}_{t,x} [V_i(\tau, X_\tau^i)] - V_i(t, X_t^i) = \mathbb{E}_{t,x} \left[\int_t^\tau \mathcal{L}^{c_s^i, w_s^i} V_i(s, X_{s-}^i) ds \right].$$

Take the supremum on both sides over the admissible controls $(c^i, w^i) \in \mathcal{A}_{t,T}^i$. Recall that the dynamic programming principle in (A.1) implies that for any stopping time τ , we have

$$\sup_{(c^i, w^i) \in \mathcal{A}_{t,\tau}^i} \mathbb{E}_{t,x} [V_i(\tau, X_\tau^i)] = V_i(t, X_t^i).$$

Therefore, we arrive at:

$$0 = \sup_{(c^i, w^i) \in \mathcal{A}_{t,\tau}^i} \mathbb{E}_{t,x} \left[\int_t^\tau \mathcal{L}^{c_s^i, w_s^i} V_i(s, X_{s-}^i) ds \right]. \quad (\text{A.6})$$

We note that this step required existence of an optimal control. For small enough δ and ϵ in (A.5), we will have $\tau = t + \delta$. Therefore, (A.6) yields

$$0 = \sup_{(c^i, w^i)} \lim_{\delta \rightarrow 0} \frac{1}{\delta} \mathbb{E}_{t,x} \left[\int_t^{t+\delta} \mathcal{L}^{c_s^i, w_s^i} V_i(s, X_{s-}^i) ds \right].$$

Finally, applying the Dominated Convergence Theorem gives

$$0 = \sup_{(c^i, w^i)} \mathcal{L}^{c_i, w_i} V_i(t, x),$$

which equals (3.2) after plugging in the definition of \mathcal{L}^{c_i, w_i} from (A.4). \square

Proof of Proposition 3.2. Both parts of this Proposition are proved nearly identically. For conciseness, full detail is only provided for case (i) where $\gamma_i = 1$.

(i): We first show that (3.2) has a separable solution. Next, the internal optimization problem is shown to be convex, and its objective function strictly concave. Finally, we show that the proposed solution is optimal.

Separability of the PDE: First we show the value function is separable. Plugging the ansatz $V_i(t, x) = g_i(t) + \log x$ into (3.2) and performing some simplification, we have:

$$0 = g_i'(t) + \sup_{c_i, w_i} \left\{ (1 - c_i)r + \sum_{j \neq i} w_{ij} \mu_j + \frac{\eta_i \mu_i}{\phi_i} + \theta_i \bar{F}_i(c_i) \log \left(\frac{x - \eta_i x}{x} \right) \sum_{j \neq i} \theta_j \bar{F}_j(c_j) \log \left(\frac{x - \phi_j w_{ij} x}{x} \right) \right\}.$$

Observe we can cancel out all remaining x 's, and obtain the following ODE for g_i :

$$0 = g_i'(t) + \frac{\eta_i \mu_i}{\phi_i} + \sup_{c_i, w_i} \left\{ (1 - c_i)r + \sum_{j \neq i} w_{ij} \mu_j + \theta_i \bar{F}_i(c_i) \log(1 - \eta_i) + \sum_{j \neq i} \theta_j \bar{F}_j(c_j) \log(1 - \phi_j w_{ij}) \right\} \quad (\text{A.7})$$

with terminal condition $g_i(T) = 0$. If \hat{c}_i and \hat{w}_{ij} are indeed the optimal solutions to the maximization in (A.7), g_i solves $g_i'(t) = -J_i^*$ with $g_i(T) = 0$, to which the solution is $g_i(t) = (T - t)J_i^*$ as desired.

Strict Concavity: Now we analyze the resulting optimization problem for c_i, w_i . Let $\mathcal{A}_i = \mathbb{R}_+ \times \prod_{j \neq i} [0, \phi_j^{-1})$ be the feasible set for this optimization problem. Clearly, \mathcal{A}_i is a convex set. We aim to solve

$$\sup_{(c_i, w_i) \in \mathcal{A}_i} (1 - c_i)r + \sum_{j \neq i} w_{ij} \mu_j + \theta_i \bar{F}_i(c_i) \log(1 - \eta_i) + \sum_{j \neq i} \theta_j \bar{F}_j(c_j) \log(1 - \phi_j w_{ij}). \quad (\text{A.8})$$

Let $h(c_i, w_i)$ denote the function to be maximized in (A.8). It is critical to observe that h is additively separable in each of its optimization variables. Therefore, we can solve for each optimal control independently. Namely, all cross-derivatives of h equal zero, which greatly simplifies the proof of strict concavity. We begin by computing partial derivatives of h with respect to each variable, which gives

$$\begin{aligned} \frac{\partial h}{\partial c_i} &= -r - \theta_i f_i(c_i) \log(1 - \eta_i) & \frac{\partial^2 h}{\partial c_i^2} &= -\theta_i f_i'(c_i) \log(1 - \eta_i) \\ \frac{\partial h}{\partial w_{ij}} &= \mu_j - \phi_j \theta_j \frac{\bar{F}_j(c_j)}{1 - \phi_j w_{ij}} & \frac{\partial^2 h}{\partial w_{ij}^2} &= -\phi_j^2 \theta_j \frac{\bar{F}_j(c_j)}{(1 - \phi_j w_{ij})^2} \quad \forall j \neq i. \end{aligned} \quad (\text{A.9})$$

Observe that within \mathcal{A}_i , we have $(1 - \phi_j w_{ij})^2 > 0$. Recall that by Assumption 1, the density function $f_j(\cdot)$ is fully supported on \mathbb{R}_+ , and $f_i'(\cdot) < 0$. Therefore, it must be the case that $\bar{F}_j(c_j) > 0$ for any admissible c_j and $\partial_{w_{ij}, w_{ij}}^2 h < 0$. Additionally, $\partial_{c_i, c_i}^2 h < 0$ because $\eta_j > 0$.

As a result, the Hessian matrix of the objective function is negative definite in the feasible region, i.e. $\nabla^2 h < 0$ everywhere in \mathcal{A}_i . Hence h is a strictly concave function; if an optimal solution to problem (A.8) exists, it is unique (Boyd and Vandenberghe 2004).

Optimality of Given Solution: To conclude, we must prove that (3.4) is optimal for bank i .

Note that $\frac{-r}{\theta_i \log(1 - \eta_i)} > 0$. Since f_i is monotonically decreasing and positive valued on \mathbb{R}_+ , its inverse $f_i^{-1}\left(\frac{-r}{\theta_i \log(1 - \eta_i)}\right)$ is well-defined if and only if $\frac{-r}{\theta_i \log(1 - \eta_i)} \leq f_i(0)$.

Since optimization problem (A.8) is convex, the first-order condition for constrained optimization is sufficient. We need only check that $y^* = (\hat{c}_i, w_i^*) \in \mathcal{A}_i$ satisfies

$$\nabla h(y^*)^T (y - y^*) \leq 0, \quad \forall y \in \mathcal{A}_i.$$

The optimization problem for h is additively separable, so this condition is equivalent to the following:

$$\begin{aligned} \partial_{c_i} h(\hat{c}_i)(c_i - \hat{c}_i) &\leq 0, \quad \forall c_i \in \mathbb{R}_+, \\ \partial_{w_{ij}} h(\hat{w}_{ij})(w_{ij} - \hat{w}_{ij}) &\leq 0, \quad \forall w_{ij} \in [0, \phi_j^{-1}), \quad \forall j \neq i. \end{aligned} \quad (\text{A.10})$$

Note that the partial derivative $\partial_{c_i} h$ in (A.9) is a function of only c_i . The same holds for the partials with respect to each w_{ij} . Note that these derivatives will depend on c_j , but this value is not controlled by bank i . Therefore, we will omit the dependence of these derivatives on the other optimization variables.

We begin with optimality of the proposed \hat{c}_i . Consider the case where $\frac{-r}{\theta_i \log(1 - \eta_i)} \leq f_i(0)$, and observe that $\partial_{c_i} h(\hat{c}_i) = 0$ using (A.9). As a result, this choice of \hat{c}_i satisfies the first-order condition for \hat{c}_i in (A.10). Conversely, let us have $\frac{-r}{\theta_i \log(1 - \eta_i)} > f_i(0)$. Since f_i is assumed to be monotone decreasing, it must be the case that $\frac{-r}{\theta_i \log(1 - \eta_i)} > \max_{c \in \mathbb{R}_+} f_i(c)$. Using again (A.9), we obtain that $\partial_{c_i} h(c) < 0$ for every $c \in \mathbb{R}_+$. In particular, we will have $\partial_{c_i} h(0) < 0$, and the first-order

condition (A.10) is satisfied by $\hat{c}_i = 0$. The proof of optimality for \hat{w}_{ij} in (3.4) follows exactly the same steps. If it is non-zero, then the proposed value solves $\partial_{w_{ij}} h(\hat{w}_{ij}) = 0$. If not, then we know that this partial derivative is negative everywhere in the feasible region for w_{ij} . Choosing $\hat{w}_{ij} = 0$ satisfies the corresponding equation in (A.10).

Concluding, we have shown that the solution given in (3.4) satisfies (A.10). Since it lies within \mathcal{A}_i , it is optimal for problem (A.8). Recall that strict concavity provides uniqueness of this solution. Finally, since all banks optimize concurrently, (3.4) is obtained by plugging the optimal value c_j^* into \hat{w}_{ij} .

(ii): The proof of this result will largely mirror that of part (i). We first check separability of the PDE. If $V_i(t, x) = g_i(t) \frac{x^{1-\gamma_i}}{1-\gamma_i}$, then we have:

$$\partial_t V_i(t, x) = g_i'(t) \frac{x^{1-\gamma_i}}{1-\gamma_i}, \quad \partial_x V_i(t, x) = g_i(t) \frac{x^{-\gamma_i}}{x}, \quad V_i(t, (1-c)x) = g_i(t) \frac{x^{1-\gamma_i}}{1-\gamma_i} (1-c)^{1-\gamma_i},$$

for all $c < 1$. Plugging these expressions into (3.2) and dividing by $x^{1-\gamma_i}$ removes any spatial variables, and we are left with the following ordinary differential equation for g_i :

$$0 = \frac{g_i'(t)}{1-\gamma_i} + g_i(t) \sup_{c_i, w_i} \left\{ (1-c_i)r + \sum_{j \neq i} w_{ij} \mu_j + \frac{\eta_i \mu_i}{\phi_i} + \theta_i \bar{F}_i(c_i) \frac{(1-\eta_i)^{1-\gamma_i} - 1}{1-\gamma_i} \right. \\ \left. + \sum_{j \neq i} \theta_j \bar{F}_j(c_j) \frac{(1-\phi_j w_{ij})^{1-\gamma_i} - 1}{1-\gamma_i} \right\},$$

with $g_i(T) = 1$.

Let \hat{c}_i and \hat{w}_{ij} be the optimal solutions to the maximization. Then we see that g_i will solve $g_i'(t) = -(1-\gamma_i)J_i^* g_i(t)$ with $g_i(T) = 1$, whose solution is $g_i(t) = \exp((1-\gamma_i)(T-t)J_i^*)$.

The optimality and uniqueness of the solution in (3.4) will be proved analogously to part (i), but by analyzing a different objective function. We are now interested in:

$$\sup_{(c_i, w_i) \in \mathcal{A}_i} (1-c_i)r + \sum_{j \neq i} w_{ij} \mu_j + \theta_i \bar{F}_i(c_i) \frac{(1-\eta_i)^{1-\gamma_i} - 1}{1-\gamma_i} + \sum_{j \neq i} \theta_j \bar{F}_j(c_j) \frac{(1-\phi_j w_{ij})^{1-\gamma_i} - 1}{1-\gamma_i}.$$

Again, this optimization problem is additively separable, which will simplify the proof of strict concavity. As before, let $h(c_i, w_i)$ denote the function to be maximized. We compute its partial derivatives to be:

$$\frac{\partial h}{\partial c_i} = -r - \theta_i f_i'(c_i) \frac{(1-\eta_i)^{1-\gamma_i} - 1}{1-\gamma_i} \quad \frac{\partial^2 h}{\partial c_i^2} = -\theta_i f_i''(c_i) \frac{(1-\eta_i)^{1-\gamma_i} - 1}{1-\gamma_i} \\ \frac{\partial h}{\partial w_{ij}} = \mu_j - \phi_j \theta_j \bar{F}_j(c_j) (1-\phi_j w_{ij})^{-\gamma_i} \quad \frac{\partial^2 h}{\partial w_{ij}^2} = \phi_j^2 \theta_j \bar{F}_j(c_j) (-\gamma_i) (1-\phi_j w_{ij})^{-\gamma_i-1}$$

Under Assumption 1, we will have both $\partial_{c_i, c_i}^2 h < 0$ and $\partial_{w_{ij}, w_{ij}}^2 h < 0$, since $w_{ij} < \phi_j^{-1}$ everywhere in \mathcal{A}_i . Therefore, h is strictly concave on \mathcal{A}_i and the optimization problem is convex. As a result, any optimal solution must be unique.

The remaining part of the proof mirrors that of part (i). Computing the gradient of h at the candidate solution in (3.4) and using the same argument will show that the first-order conditions in (A.10) are satisfied. Since this point is feasible, it must be optimal. \square

Proof of Corollary 3.3. We proceed with a standard verification argument. We need to show that if ψ is a solution to the PDE (3.2) and it is $\mathcal{C}^{1,1}$ ($[0, T], \mathbb{R}_+$), then it is equal to the value function. Since the proposed solutions solve the PDE and they are indeed $\mathcal{C}^{1,1}$, this will conclude.

Fix $t < T$, and choose $\{\hat{c}_s^i, \hat{w}_s^i\}_{s \in [t, T]}$ be some admissible controls. We apply Itô's formula to $\psi(s, X_s^i)$ between t and some stopping time τ^n – to be chosen optimally later. This yields, using the notation introduced in the proof of Proposition 3.1, the following:

$$\begin{aligned} \psi(\tau^n, X_{\tau^n}^i) &= \psi(t, X_t^i) + \int_t^{\tau^n} \mathcal{L}^{\hat{c}_s^i, \hat{w}_s^i} \psi(s, X_s^i) ds + \int_t^{\tau^n} [\psi(s, X_{s-}^i - \eta_i X_{s-}^i) - \psi(s, X_{s-}^i)] dM_s^i \\ &\quad + \sum_{j \neq i} \int_t^{\tau^n} [\psi(s, X_{s-}^i - \phi_j \hat{w}_s^{ij} X_{s-}^i) - \psi(s, X_{s-}^i)] dM_s^j. \end{aligned}$$

Recall that the compensated jump process $\{M_t^k\}_{t \geq 0}$ is a martingale. Taking the expectation conditioned on $X_t^i = x$, we obtain:

$$\begin{aligned} \mathbb{E}_{t,x} [\psi(\tau^n, X_{\tau^n}^i)] &= \psi(t, x) + \mathbb{E}_{t,x} \left[\int_t^{\tau^n} \mathcal{L}^{\hat{c}_s^i, \hat{w}_s^i} \psi(s, X_s^i) ds \right] \\ &\quad + \mathbb{E}_{t,x} \left[\int_t^{\tau^n} [\psi(s, X_{s-}^i - \eta_i X_{s-}^i) - \psi(s, X_{s-}^i)] dM_s^i \right] \\ &\quad + \sum_{j \neq i} \mathbb{E}_{t,x} \left[\int_t^{\tau^n} [\psi(s, X_{s-}^i - \phi_j \hat{w}_s^{ij} X_{s-}^i) - \psi(s, X_{s-}^i)] dM_s^j \right]. \end{aligned}$$

If we choose $\tau^n = (T - \frac{1}{n}) \wedge \inf \{s \in [t, T], X_s^i \leq \frac{1}{n} \text{ or } X_s^i \geq n\}$, then for every n the expectation of each stochastic integral is zero and we have:

$$\mathbb{E}_{t,x} [\psi(\tau^n, X_{\tau^n}^i)] = \psi(t, x) + \mathbb{E}_{t,x} \left[\int_t^{\tau^n} \mathcal{L}^{\hat{c}_s^i, \hat{w}_s^i} \psi(s, X_s^i) ds \right].$$

Taking the limit as $n \rightarrow \infty$, we will have $\tau^n \rightarrow T$. Furthermore, since ψ satisfies the terminal condition (by assumption) and everything is bounded, an application of dominated convergence yields:

$$\mathbb{E}_{t,x} [U_i(X_T^i)] = \psi(t, x) + \mathbb{E}_{t,x} \left[\int_t^T \mathcal{L}^{\hat{c}_s^i, \hat{w}_s^i} \psi(s, X_s^i) ds \right]. \quad (\text{A.11})$$

First, we choose the controls in (A.11) to be given by the optimal solution of Proposition 3.2. Then, we will have $\mathcal{L}^{\hat{c}_s^i, \hat{w}_s^i} \psi(s, X_s^i) = 0$ for all $s \in [t, \tau^n]$, and consequentially:

$$\psi(t, x) = \mathbb{E}_{t,x} [U_i(X_T^i)].$$

Note that only the terminal wealth X_T^i in the right-hand side depends on the controls $(\hat{c}_s^i, \hat{w}_s^i)$. After taking the supremum we obtain

$$\psi(t, x) \leq \sup_{\{\hat{c}_s^i, \hat{w}_s^i\}_{s \in [t, T]}} \mathbb{E}_{t,x} [U_i(X_T^i)] = V_i(t, x). \quad (\text{A.12})$$

Next, we fix any control $(\hat{c}_s^i, \hat{w}_s^i)$. Then, in (A.11) we will have $\mathcal{L}^{\hat{c}_s^i, \hat{w}_s^i} \psi(s, X_s^i) \leq 0$, and the result is:

$$\psi(t, x) \geq \mathbb{E}_{t,x} [U_i(X_T^i)].$$

Note again that only X_T^i depends on the controls. However, since this inequality holds for any admissible control we can take the supremum over both sides to give

$$\psi(t, x) \geq \sup_{\{c_s^i, w_s^i\}_{s \in [t, T]}} \mathbb{E}_{t, x} [U_i(X_T^i)] = V_i(t, x). \quad (\text{A.13})$$

Combining (A.12) and (A.13) shows that $\psi = V_i$. This implies that the optimal values to the maximization problem in the PDE for ψ are indeed the optimal controls. Since the explicit solutions given by Proposition 3.2 are once continuously differentiable in both time and space, then they are equal to the value function. \square

A.2 Centralized Network

Proof of Proposition 3.4. This proof is only a minor adaptation of the proof of Proposition 3.1. First, the application of Itô's formula to the value function $V(t, X_t^1, \dots, X_t^n)$ yields more terms, but remains simple as the jump processes are mutually independent. Namely, the generator is given by

$$\begin{aligned} \mathcal{L}^{c, w} \psi = & \partial_t \psi + \sum_{i=1}^n \left(\left[(1 - c_i)r + \sum_{j \neq i} w_{ij} \mu_j + \frac{\eta_i \mu_i}{\phi_i} \right] x_i \partial_{x_i} \psi \right. \\ & \left. + \theta_i \bar{F}_i(c_i) \left[\psi(t, x_1(1 - \phi_i w_{1i}), \dots, x_i(1 - \eta_i), \dots, x_n(1 - \phi_i w_{ni})) - \psi \right] \right), \end{aligned}$$

where ψ is evaluated at (t, x_1, \dots, x_n) where unspecified.

Next, to apply dominated convergence, the choice of the stopping time τ must ensure that all stopped processes $X_\tau^1, \dots, X_\tau^n$ are bounded away from zero. We can therefore choose:

$$\tau = (t + \delta) \wedge \min_i \left\{ \inf \left\{ s \in [t, T], X_s^i \leq \epsilon \text{ or } X_s^i \geq \frac{1}{\epsilon} \right\} \right\},$$

and conclude as in the previous result. \square

Proof of Proposition 3.5. The outline of this proof is similar to that of Prop. 3.2, but with greater complexity, and hence requiring additional assumptions to establish our results. We begin by discussing each of these. First, logarithmic utility functions are needed so that (3.7) admits a separable solution. We note that if the planner sought to maximize the product of banks' utilities, it would be necessary to assume that $\gamma_i \neq 1$ for all i . This assumption is used for existence of a separable solution to (3.7).

The first condition in Assumption 2 concerns the shock densities f_i . In particular, (3.8) is satisfied by the family of exponential distributions ($f_i(x) = \lambda_i^{-1} e^{-\frac{x}{\lambda_i}}$, for some parameter $\lambda_i > 0$) and power distributions ($f_i(x) = \frac{(\alpha_i^{-1} - 1)x_0^{\alpha_i^{-1} - 1}}{(x + x_0)^{\alpha_i^{-1}}}$, for any $x_0 > 0$ and $\alpha_i < 1$). We note that this condition is not necessary for uniqueness, but is used for establishing monotonicity of a first-order condition for optimality by bounding the second derivative with an exponentially decaying function.

Finally, the inequalities on $\Gamma(\eta_i; 1)$ will ensure that either (i): strict concavity of the objective function holds, or (ii) there exists only a single solution to the necessary first-order conditions. However, these inequalities do not rule out the possibility of a corner solution of $c_i^* = 0$ or $w_i^* = 0$ – as shown in (3.10). Of particular interest, the optimal decentralized and centralized allocations for c_i and w_i will coincide whenever either $c_i^* = 0$ or $w_i^* = 0$ in the planner's optimum.

Separability of PDE and Maximization: Recall that the PDE for the value function derived in Proposition 3.4 is:

$$0 = \partial_t V + \sup_{c, w} \left\{ \sum_{i=1}^n \left(\left[(1 - c_i) r + \sum_{j \neq i} w_{ij} \mu_j + \frac{\eta_i \mu_i}{\phi_i} \right] x_i \partial_{x_i} V + \theta_i \bar{F}_i(c_i) \left[V(t, x_1(1 - \phi_i w_{1i}), \dots, x_i(1 - \eta_i), \dots, x_n(1 - \phi_i w_{ni})) - V \right] \right) \right\} \quad (\text{A.14})$$

$$V(T, x_1, \dots, x_n) = \sum_{i=1}^n U_i(x_i).$$

By assumption, each bank's utility function is given by $U_i(x_i) = \log x_i$, i.e. $\gamma_i = 1$ for all i . Consider the following ansatz: $V(t, x_1, \dots, x_n) = g(t) + \sum_i \log x_i$. Substituting into (A.14), we obtain:

$$0 = g'(t) + \sup_{c, w} \sum_{i=1}^n \left((1 - c_i) r + \sum_{j \neq i} w_{ij} \mu_j + \frac{\eta_i \mu_i}{\phi_i} - \theta_i \bar{F}_i(c_i) \left[\Gamma(\eta_i; 1) + \sum_{j \neq i} \Gamma(\phi_i w_{ji}; 1) \right] \right) \quad (\text{A.15})$$

with $g(T) = 0$. The spatial variables will cancel and we are left with an ordinary differential equation for g . We now rewrite the following sum:

$$\sum_{i=1}^n \sum_{j \neq i} w_{ij} \mu_j = \sum_{i=1}^n \sum_{j \neq i} w_{ji} \mu_i.$$

Observe that for $k, j \neq i$, we will have $w_{ji} = w_{ki}$. That is, all $j \neq i$ banks will invest the same fraction of their wealth in bank i 's project. This can be seen in two ways. First, in the decentralized setting, the amount w_{ji} depended on bank j only through their risk aversion coefficient γ_j . Since in this Proposition we have assumed that $\gamma_i = 1$ for all i , the result follows. This can also be seen by computing the first-order conditions in (A.15) for w_{ji} and w_{ki} , and noticing that they are identical. Let this fraction be denoted by w_i . This allows us to further simplify (A.15) and obtain

$$0 = g'(t) + \sum_{i=1}^n \frac{\eta_i \mu_i}{\phi_i} + \sup_{c, w} \sum_{i=1}^n \left((1 - c_i) r + (n - 1) w_i \mu_i - \theta_i \bar{F}_i(c_i) \left[\Gamma(\eta_i; 1) + (n - 1) \Gamma(\phi_i w_i; 1) \right] \right).$$

This maximization is additively separable between each pair (c_i, w_i) , indexed by i . Let $\mathcal{A}_i = \mathbb{R}_+ \times [0, \phi_i^{-1})$ denote the admissible values for (c_i, w_i) . Then, the optimal allocation is found by solving:

$$\sum_{i=1}^n \sup_{(c_i, w_i) \in \mathcal{A}_i} h_i(c_i, w_i), \quad (\text{A.16})$$

where $h_i(c_i, w_i) = -rc_i + (n - 1)\mu_i w_i - \theta_i \bar{F}_i(c_i) \left[\Gamma(\eta_i; 1) + (n - 1)\Gamma(\phi_i w_i; 1) \right]$ for each i .

Reduction to Univariate Optimization: We first maximize over w_i and then c_i given the optimal w_i . Given a value of c_i , we seek to find the optimal value of w_i . We can compute

$$\begin{aligned} \frac{\partial h_i}{\partial w_i}(c_i, w_i) &= (n - 1)\mu_i - (n - 1) \frac{\phi_i \theta_i \bar{F}_i(c_i)}{1 - \phi_i w_i} \\ \frac{\partial^2 h_i}{\partial w_i^2}(c_i, w_i) &= -(n - 1) \frac{\phi_i^2 \theta_i \bar{F}_i(c_i)}{(1 - \phi_i w_i)^2}. \end{aligned} \quad (\text{A.17})$$

Notice that the second derivative in this expression is always strictly negative. Hence, given c_i , the optimization problem over $w_{\cdot i}$ is strictly concave. This implies that the first-order conditions are sufficient, and that any optimal solution is unique. Let $w_{\cdot i}^*(c_i)$ denote the optimal solution given c_i . It must satisfy the following necessary first-order condition:

$$\frac{\partial h_i}{\partial w_{\cdot i}}(c_i, w_{\cdot i}^*(c_i))(w_{\cdot i} - w_{\cdot i}^*(c_i)) \leq 0, \quad \forall w_{\cdot i} \in [0, \phi_i^{-1}].$$

Using (A.17), it is easy to check that this condition is satisfied by the following:

$$w_{\cdot i}^*(c_i) = \begin{cases} \frac{1}{\phi_i} \left(1 - \frac{\phi_i \theta_i \bar{F}_i(c_i)}{\mu_i}\right) & \text{if } \frac{\phi_i \theta_i \bar{F}_i(c_i)}{\mu_i} \leq 1 \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.18})$$

This value is uniquely defined, and exists for any choice of c_i . We then rewrite each maximization in (A.16) as

$$\sup_{(c_i, w_{\cdot i}) \in \mathcal{A}_i} h_i(c_i, w_{\cdot i}) = \sup_{c_i \geq 0} h_i^*(c_i), \quad (\text{A.19})$$

where $h_i^*(c_i) = h_i(c_i, w_{\cdot i}^*(c_i))$.

Existence of an Optimal Solution: We now prove existence of an optimal solution to (A.19). Observe that for large enough c_i , we will have $w_{\cdot i}^*(c_i) = \frac{1}{\phi_i} \left(1 - \frac{\phi_i \theta_i \bar{F}_i(c_i)}{\mu_i}\right)$. For such c_i we obtain

$$h_i^*(c_i) = -rc_i + (n-1)\mu_i \left[\frac{1}{\phi_i} \left(1 - \frac{\phi_i \theta_i \bar{F}_i(c_i)}{\mu_i}\right) \right] - \theta_i \bar{F}_i(c_i) \left[\Gamma(\eta_i; 1) - (n-1) \log \left(\frac{\phi_i \theta_i \bar{F}_i(c_i)}{\mu_i} \right) \right].$$

As $c_i \rightarrow \infty$, we will have $\bar{F}_i(c_i) \rightarrow 0$. Since we can write

$$\bar{F}_i(c_i) \log \left(\frac{\phi_i \theta_i \bar{F}_i(c_i)}{\mu_i} \right) = \bar{F}_i(c_i) \left[\log \left(\frac{\phi_i \theta_i}{\mu_i} \right) + \log \bar{F}_i(c_i) \right],$$

and $x \log x \xrightarrow{x \rightarrow 0} 0$, we will have $\lim_{c_i \rightarrow \infty} h_i^*(c_i) = -\infty$.

This limit is sufficient for existence of an optimal solution to (A.19). Fix some $K < 0$. Since we have shown $h_i^*(c_i) \xrightarrow{c_i \rightarrow \infty} -\infty$, we know that $\exists C \in \mathbb{R}_+ : h_i^*(c_i) < K, \forall c_i > C$. By continuity of h_i^* , the set $\mathcal{B} = \{c_i \in \mathbb{R}_+ : h_i^*(c_i) \geq K\}$ is compact. We can conclude by the Extreme Value Theorem that there exists a globally optimal value of h_i^* within \mathcal{B} . Moreover, as long as \mathcal{B} is non-empty, any point in \mathcal{B} achieves higher objective value than any point in its complement. By taking K to be a large enough negative number, we can ensure that $\mathcal{B} \neq \emptyset$.

System of Equations for Optimum: The expression (A.18) gives us the second equation in the system (3.10). For the other equation, we must analyze the first-order condition for c_i in (A.16). Taking derivatives with respect to c_i , we obtain

$$\begin{aligned} \frac{\partial h_i}{\partial c_i}(c_i, w_{\cdot i}) &= -r + \theta_i f_i(c_i) \left[\Gamma(\eta_i; 1) + (n-1) \Gamma(\phi_i w_{\cdot i}; 1) \right] \\ \frac{\partial^2 h_i}{\partial c_i^2}(c_i, w_{\cdot i}) &= -\theta_i f_i'(c_i) \left[\Gamma(\eta_i; 1) + (n-1) \Gamma(\phi_i w_{\cdot i}; 1) \right]. \end{aligned} \quad (\text{A.20})$$

Notice that the second derivative is also negative everywhere – although this does not imply that the objective function h_i is concave. We proceed similarly as before, seeking to define an optimal value of c_i for any given $w_{\cdot i}$. Let this be denoted $c_i^*(w_{\cdot i})$. It must satisfy:

$$\frac{\partial h_i}{\partial c_i}(c_i^*(w_{\cdot i}), w_{\cdot i})(c_i - c_i^*(w_{\cdot i})) \leq 0, \quad \forall c_i \in \mathbb{R}_+.$$

Using (A.20), we can see that this will be satisfied whenever

$$c_i^*(w_{\cdot i}) = \begin{cases} f_i^{-1} \left(\frac{r}{\theta_i[\Gamma(\eta_i; 1) + (n-1)\Gamma(\phi_i w_{\cdot i}; 1)]} \right) & \text{if } f_i(0) \leq \frac{r}{\theta_i[\Gamma(\eta_i; 1) + (n-1)\Gamma(\phi_i w_{\cdot i}; 1)]} \\ 0 & \text{otherwise.} \end{cases}$$

With (A.18), we obtain the system (3.10).

Uniqueness: It remains only to show that the optimal solution to (A.19) is unique. We return to our analysis of the univariate optimization problem in (A.19). The necessary first-order condition for optimality of c_i^* is

$$\frac{dh_i^*}{dc_i}(c_i^*)(c_i - c_i^*) \leq 0, \quad \forall c_i \in \mathbb{R}_+. \quad (\text{A.21})$$

We proceed by showing that there exists only a single c_i^* satisfying this expression, and since existence has been proved, it must be the optimal solution. Recall that $\tilde{c}_i = F_i^{-1} \left(\left[1 - \frac{\mu_i}{\phi_i \theta_i} \right]_+ \right)$, and we have $w_{\cdot i}^*(c_i) = 0$ if and only if $c_i \leq \tilde{c}_i$.

The reduced objective function $h_i^*(c_i)$, after substituting in (A.18), can be written as:

$$h_i^*(c_i) = -rc_i - \theta_i \bar{F}_i(c_i) \Gamma(\eta_i; 1) + \begin{cases} (n-1) \left[\frac{\mu_i}{\phi_i} + \theta_i \bar{F}_i(c_i) \left(\log \left(\frac{\phi_i \theta_i \bar{F}_i(c_i)}{\mu_i} \right) - 1 \right) \right] & \text{if } c_i \geq \tilde{c}_i \\ 0 & \text{otherwise.} \end{cases}$$

Taking the derivative with respect to c_i , we obtain

$$\frac{dh_i^*}{dc_i}(c_i) = -r + \theta_i f_i(c_i) \Gamma(\eta_i; 1) - \begin{cases} \theta_i f_i(c_i) (n-1) \log \left(\frac{\phi_i \theta_i \bar{F}_i(c_i)}{\mu_i} \right) & \text{if } c_i \geq \tilde{c}_i \\ 0 & \text{otherwise,} \end{cases}$$

and the second derivative equals

$$\frac{d^2 h_i^*}{dc_i^2}(c_i) = \theta_i f_i'(c_i) \Gamma(\eta_i; 1) + \theta_i (n-1) \begin{cases} \frac{f_i(c_i)^2}{1 - F_i(c_i)} - f_i'(c_i) \log \left(\frac{\phi_i \theta_i \bar{F}_i(c_i)}{\mu_i} \right) & \text{if } c_i \geq \tilde{c}_i \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.22})$$

Note that we are evaluating the right derivatives at $c_i = \tilde{c}_i$, where this function is not differentiable.

In the regime $c_i < \tilde{c}_i$, we will always have $\frac{d^2 h_i^*}{dc_i^2}(c_i) < 0$. If this were also true for $c_i \geq \tilde{c}_i$, then the objective function would be strictly concave, and uniqueness would follow. We now prove that if $\frac{d^2 h_i^*}{dc_i^2}(x) < 0$, then $h_i^*(\cdot)$ is strictly concave on $[x, \infty)$. In particular, by plugging in $x = \tilde{c}_i$ we conclude uniqueness of the optimum.

Let us compute an additional derivative of $h_i^*(\cdot)$:

$$\frac{d^3 h_i^*}{dc_i^3}(c_i) = \theta_i f_i''(c_i) \Gamma(\eta_i; 1) + \theta_i (n-1) \begin{cases} \frac{f_i(c_i)^2}{F_i(c_i)} \left[\frac{f_i(c_i)}{F_i(c_i)} + 3 \frac{f_i'(c_i)}{f_i(c_i)} \right] - f_i''(c_i) \log \left(\frac{\phi_i \theta_i \bar{F}_i(c_i)}{\mu_i} \right) & \text{if } c_i \geq \tilde{c}_i \\ 0 & \text{otherwise.} \end{cases}$$

Observe that when we have $c_i \geq \tilde{c}_i$, a bit of algebra yields

$$\frac{d^3 h_i^*}{dc_i^3}(c_i) = \frac{f_i''(c_i)}{f_i'(c_i)} \frac{d^2 h_i^*}{dc_i^2}(c_i) + \frac{(n-1)\theta_i f_i(c_i)^2}{F_i(c_i)} \left[\frac{f_i(c_i)}{F_i(c_i)} + 3 \frac{f_i'(c_i)}{f_i(c_i)} - \frac{f_i''(c_i)}{f_i'(c_i)} \right].$$

If, as assumed in this Proposition, we have $\frac{f_i(c_i)}{F_i(c_i)} + 3 \frac{f_i'(c_i)}{f_i(c_i)} - \frac{f_i''(c_i)}{f_i'(c_i)} < 0$ for all $c_i \geq 0$, then it will follow that

$$\frac{d^3 h_i^*}{dc_i^3}(c_i) < \frac{f_i''(c_i)}{f_i'(c_i)} \frac{d^2 h_i^*}{dc_i^2}(c_i).$$

Applying Grönwall's inequality, we see that

$$\frac{d^2 h_i^*}{dc_i^2}(b) < \frac{d^2 h_i^*}{dc_i^2}(a) \exp \left(\int_a^b \frac{f_i''(s)}{f_i'(s)} ds \right),$$

for any $\tilde{c}_i \leq a < b$. As a consequence, if $\frac{d^2 h_i^*}{dc_i^2}(a) \leq 0$, then $\frac{d^2 h_i^*}{dc_i^2}(b) < 0$ for all $b > a$.

Rewriting (A.22), we obtain:

$$\frac{d^2 h_i^*}{dc_i^2}(\tilde{c}_i) = \begin{cases} \theta_i f_i'(0) \left[\Gamma(\eta_i; 1) - (n-1) \log \left(\frac{\phi_i \theta_i}{\mu_i} \right) \right] + \theta_i (n-1) f_i(0)^2 & \text{if } \tilde{c}_i = 0 \\ \theta_i f_i'(\tilde{c}_i) \Gamma(\eta_i; 1) + \theta_i (n-1) \frac{\phi_i \theta_i f_i(\tilde{c}_i)^2}{\mu_i} & \text{otherwise.} \end{cases}$$

For i satisfying

$$\Gamma(\eta_i; 1) > \begin{cases} (n-1) \left[\log \left(\frac{\phi_i \theta_i}{\mu_i} \right) - \frac{f_i(0)^2}{f_i'(0)} \right] & \text{if } \tilde{c}_i = 0 \\ -(n-1) \frac{\phi_i \theta_i f_i(\tilde{c}_i)^2}{\mu_i f_i'(\tilde{c}_i)} & \text{otherwise.} \end{cases} \quad (\text{A.23})$$

in the assumption (3.9), we see that $\frac{d^2 h_i^*}{dc_i^2}(\tilde{c}_i) < 0$. By our application of Grönwall's inequality, we can conclude that h_i^* must be strictly concave, and hence the optimum is unique.

Now, we turn to the banks i satisfying

$$\Gamma(\eta_i; 1) > \begin{cases} \frac{r}{\theta_i f_i(0)} + (n-1) \log \left(\frac{\phi_i \theta_i}{\mu_i} \right) & \text{if } \tilde{c}_i = 0 \\ \frac{r}{\theta_i f_i(\tilde{c}_i)} & \text{otherwise.} \end{cases} \quad (\text{A.24})$$

We can compute:

$$\frac{dh_i^*}{dc_i}(\tilde{c}_i) = -r + \begin{cases} \theta_i f_i(0) \left[\Gamma(\eta_i; 1) - (n-1) \log \left(\frac{\phi_i \theta_i}{\mu_i} \right) \right] & \text{if } \tilde{c}_i = 0 \\ \theta_i f_i(\tilde{c}_i) \Gamma(\eta_i; 1) & \text{otherwise.} \end{cases}$$

By (A.24), we have $\frac{dh_i^*}{dc_i}(\tilde{c}_i) > 0$. Since $\frac{d^2 h_i^*}{dc_i^2}(c_i) < 0$ for all $c_i < \tilde{c}_i$, we cannot have any points satisfying the first-order condition (A.21) in $[0, \tilde{c}_i]$. However, we do know that there must exist an optimal solution, so therefore it must lie within (\tilde{c}_i, ∞) . At such a point c_i^* , we must have $\frac{dh_i^*}{dc_i}(c_i^*) = 0$, and also $\frac{d^2 h_i^*}{dc_i^2}(c_i^*) \leq 0$, which are the two necessary conditions for optimality of c_i^* when it lies in the interior of the feasible region. By the same conclusion using Grönwall's inequality, we must have $\frac{d^2 h_i^*}{dc_i^2}(c_i) < 0$, and hence $\frac{dh_i^*}{dc_i}(c_i) < 0$ for any $c_i > c_i^*$. Hence, only this choice of c_i^* will satisfy the necessary first-order conditions, and as a result it must be unique.

Since we require all i to satisfy at least one of (A.24) or (A.23), the optimal solutions to each of the n optimization problems in (A.16) must be unique. \square

Proof of Corollary 3.6. The proof of this result mirrors the proof of Corollary 3.3, and therefore we omit many details.

Fix some time $t < T$, at which we have $X_t^i = x_i$. We again choose some admissible controls $\{c_s^i, w_s^i\}_{s \in [t, T]}$. We then apply Itô's formula, which only differs in yielding a few more terms. Namely, we will need to use the generator defined in Section (A.2), and the stochastic integrands will be slightly more complex. Next, to apply dominated convergence, our choice of the stopping time τ^n must ensure that each of the wealth processes $\{X_s^1\}_{s \geq 0} \cdots \{X_s^n\}_{s \geq 0}$, is bounded at time τ^n . Therefore, we choose

$$\tau^n = \left(T - \frac{1}{n}\right) \wedge \min_i \left\{ \inf \{s \geq t, |X_s^i - X_t^i| \geq n\} \right\}$$

and conclude identically. \square

A.3 Differences in Optima

Lemma A.1. *Assume that $w_i^* > 0$. If the shock density satisfies: $f_i(x) \geq \kappa_{i,L} e^{-\frac{x}{\lambda_{i,L}}}$, for all x and fixed constants $\lambda_{i,L} > 0$ and $\kappa_{i,L} > 0$, then*

$$c_i^* \geq \lambda_{i,L} \log \left(\frac{\theta_i \kappa_{i,L} \Gamma(\phi_i \hat{w}_i; 1)}{r} \right) + \lambda_{i,L} \log(n-1).$$

In particular, the planner's optimal cash holdings asymptotically grow at least logarithmically in n . Furthermore, if for all x we also have:

$$f_i(x) \leq \kappa_{i,U} e^{-\frac{x}{\lambda_{i,U}}}$$

for $\lambda_{i,L} \leq \lambda_{i,U}$ and $\kappa_{i,L} \leq \kappa_{i,U}$, then

(i) **Upper Bound:**

$$c_i^* \leq \lambda_{i,U} \log \left(\frac{\theta_i \kappa_{i,U} C_U}{r} \right) + \lambda_{i,U} \log((n-1) \log(n)),$$

where $C_U > 3$ depends on all model parameters (including $\lambda_{i,L}$ and $\lambda_{i,U}$), but does not explicitly grow with n . As a result, $\lim_{n \rightarrow \infty} \frac{c_i^}{\log(n)} \leq \lambda_{i,U}$.*

(ii) **Lower Bound:**

$$c_i^* \geq \lambda_{i,L} \log \left(\frac{\theta_i \kappa_{i,L} \lambda_{i,L}}{r \lambda_{i,U}} \right) + \lambda_{i,L} \log \left((n-1) \left[\log(n-1) - \frac{\lambda_{i,U}}{\lambda_{i,L}} \log(C_L) \right] \right),$$

for $C_L > 0$ depending only on i 's parameters. Hence, $\lim_{n \rightarrow \infty} \frac{c_i^}{\log(n)} \geq \lambda_{i,L}$.*

Combining the two limiting bounds, we have $c_i^ = \Theta(\log(n))$.*

The proof follows from iterating through upper (and lower) bounds for c_i^* using the system of equations in (3.10), and beginning from crude estimates. It is possible to use the same techniques in this proof to obtain bounds when the density has power-law tails. While the results are not qualitatively different, we are unable to achieve the tight bound that appears in Proposition 4.1 when the shock distribution is itself a power-law. The main result can be seen in Appendix B.

Proof of Lemma A.1. The main idea in this proof is to first establish crude bounds of:

$$\hat{c}_i \leq c_i^* \leq Kn^2,$$

for a suitable choice of K . This then allows us to improve the bounds on c_i^* itself through the relationship

$$c_i^* = f_i^{-1} \left(\frac{r}{\theta_i \left[\Gamma(\eta_i; 1) - (n-1) \log \left(\frac{\phi_i \theta_i \bar{F}_i(c_i^*)}{\mu_i} \right) \right]} \right),$$

using the assumptions of a super- and sub-exponential density.

Through a direct computation with the explicit solutions in Propositions 3.2 and 3.5, we can write

$$\begin{aligned} V(t, x_1, \dots, x_n) - \sum_{i=1}^n V_i(t, x_i) &= (T-t) \left[J_C^* - \sum_{i=1}^n J_i^* \right] \\ &= (T-t) \sum_{i=1}^n \left[-r(c_i^* - \hat{c}_i) + (n-1)\mu_i(w_{\cdot i}^* - \hat{w}_{\cdot i}) \right. \\ &\quad \left. - \theta_i \bar{F}_i(c_i^*) [\Gamma(\eta_i; 1) + (n-1)\Gamma(\phi_i w_{\cdot i}^*)] \right. \\ &\quad \left. + \theta_i \bar{F}_i(\hat{c}_i) [\Gamma(\eta_i; 1) + (n-1)\Gamma(\phi_i \hat{w}_{\cdot i})] \right] \end{aligned}$$

Observe that using the definitions, we have $w_{\cdot i}^* - \hat{w}_{\cdot i} = \frac{\theta_i}{\mu_i} (\bar{F}_i(\hat{c}_i) - \bar{F}_i(c_i^*))$. Plugging this expression in and rearranging terms, we obtain:

$$\begin{aligned} \frac{g(t) - \sum_{i=1}^n g_i(t)}{T-t} &= \sum_{i=1}^n \left[-r(c_i^* - \hat{c}_i) + \theta_i (\bar{F}_i(\hat{c}_i) - \bar{F}_i(c_i^*)) [(n-1) + \Gamma(\eta_i; 1)] \right. \\ &\quad \left. + \theta_i (n-1) [\bar{F}_i(\hat{c}_i) \Gamma(\phi_i \hat{w}_{\cdot i}; 1) - \bar{F}_i(c_i^*) \Gamma(\phi_i w_{\cdot i}^*; 1)] \right]. \end{aligned} \tag{A.25}$$

Since we know the gap in (A.25) must be positive, we can write:

$$\begin{aligned} \sum_{i=1}^n r c_i^* &\leq \sum_{i=1}^n \left[r \hat{c}_i + \theta_i (\bar{F}_i(\hat{c}_i) - \bar{F}_i(c_i^*)) [(n-1) + \Gamma(\eta_i; 1)] \right. \\ &\quad \left. + \theta_i (n-1) [\bar{F}_i(\hat{c}_i) \Gamma(\phi_i \hat{w}_{\cdot i}; 1) - \bar{F}_i(c_i^*) \Gamma(\phi_i w_{\cdot i}^*; 1)] \right] \\ &\leq \sum_{i=1}^n \left[r \hat{c}_i + \theta_i \bar{F}_i(\hat{c}_i) [(n-1) + \Gamma(\eta_i; 1)] \right. \\ &\quad \left. + \theta_i (n-1) [\bar{F}_i(\hat{c}_i) \Gamma(\phi_i \hat{w}_{\cdot i}; 1)] \right], \end{aligned}$$

which follows by dropping the final term and since $\bar{F}_i(c_i^*) \geq 0$. A crude bound implies that

$$\begin{aligned} r c_i^* &\leq \sum_{i=1}^n (n-1) \left[r \hat{c}_i + \theta_i \bar{F}_i(\hat{c}_i) [1 + \Gamma(\eta_i; 1) + \Gamma(\phi_i \hat{w}_{\cdot i}; 1)] \right] \\ c_i^* &\leq Kn^2, \end{aligned}$$

where $K = \max_i \left\{ \hat{c}_i + \frac{\theta_i}{r} \bar{F}_i(\hat{c}_i) [1 + \Gamma(\eta_i; 1) + \Gamma(\phi_i \hat{w}_i; 1)] \right\}$ does not depend explicitly on n . Since $w_i^* \geq 0$, it is also easy to see that $c_i^* \geq \hat{c}_i$. Both these bounds will be useful starting points for the proof.

- (i) **Upper Bound:** We first prove the upper bound for c_i^* . First, since $f_i(x) \leq \kappa_{i,U} e^{-\frac{x}{\lambda_{i,U}}}$ and both functions are decreasing, we will have $f_i^{-1}(y) \leq \lambda_{i,U} \log\left(\frac{\kappa_{i,U}}{y}\right)$, and it follows from the system of equations (3.10) that

$$c_i^* \leq \lambda_{i,U} \log \left(\frac{\theta_i \kappa_{i,U} \left[\Gamma(\eta_i; 1) - (n-1) \log \left(\frac{\phi_i \theta_i \bar{F}_i(c_i^*)}{\mu_i} \right) \right]}{r} \right).$$

Now, using $f_i(x) \geq \kappa_{i,L} e^{-\frac{x}{\lambda_{i,L}}}$, we know that $\bar{F}_i(c_i^*) = \int_{c_i^*}^{\infty} f_i(u) du \geq \kappa_{i,L} \lambda_{i,L} e^{-\frac{c_i^*}{\lambda_{i,L}}}$, and write:

$$\begin{aligned} c_i^* &\leq \lambda_{i,U} \log \left(\frac{\theta_i \kappa_{i,U} \left[\Gamma(\eta_i; 1) - (n-1) \log \left(\frac{\phi_i \theta_i \kappa_{i,L} \lambda_{i,L}}{\mu_i} \right) + (n-1) \frac{c_i^*}{\lambda_{i,L}} \right]}{r} \right) \\ &\leq \lambda_{i,U} \log \left(\frac{\theta_i \kappa_{i,U} \left[\Gamma(\eta_i; 1) - (n-1) \log \left(\frac{\phi_i \theta_i \kappa_{i,L} \lambda_{i,L}}{\mu_i} \wedge 1 \right) + (n-1) \frac{c_i^*}{\lambda_{i,L}} \right]}{r} \right). \end{aligned} \quad (\text{A.26})$$

Since each of the three terms in the brackets is non-negative, we can upper bound this quantity by:

$$c_i^* \leq \lambda_{i,U} \log \left(\frac{\theta_i \kappa_{i,U} \left[\Gamma(\eta_i; 1) - \log \left(\frac{\phi_i \theta_i \kappa_{i,L} \lambda_{i,L}}{\mu_i} \wedge 1 \right) + \lambda_{i,L}^{-1} \right] n c_i^*}{r} \right),$$

and we define $D = \Gamma(\eta_i; 1) - \log \left(\frac{\phi_i \theta_i \kappa_{i,L} \lambda_{i,L}}{\mu_i} \wedge 1 \right) + \lambda_{i,L}^{-1}$ for convenience. Recall that we obtained a crude upper bound of $c_i^* \leq K n^2$, which, when plugged in, yields:

$$c_i^* \leq \lambda_{i,U} \log \left(\frac{\theta_i \kappa_{i,U} D K n^3}{r} \right).$$

This is a significantly tighter bound than $K n^2$. Therefore, we plug it back into (A.26). By simplifying and bounding the term in the logarithm, we compute:

$$\begin{aligned} \frac{c_i^*}{\lambda_{i,U}} &\leq \log \left(\frac{\theta_i \kappa_{i,U} \left[\Gamma(\eta_i; 1) - (n-1) \log \left(\frac{\phi_i \theta_i \kappa_{i,L} \lambda_{i,L}}{\mu_i} \right) + (n-1) \frac{\lambda_{i,U}}{\lambda_{i,L}} \log \left(\frac{\theta_i \kappa_{i,U} D K n^3}{r} \right) \right]}{r} \right) \\ &\leq \log \left(\frac{\theta_i \kappa_{i,U} \left[\Gamma(\eta_i; 1) + (n-1) \left[\log \left(\frac{\mu_i}{\phi_i \theta_i \kappa_{i,L} \lambda_{i,L}} \left(\frac{\theta_i \kappa_{i,U} D K}{r} \right)^{\frac{\lambda_{i,U}}{\lambda_{i,L}}} \vee 1 \right) + 3 \frac{\lambda_{i,U}}{\lambda_{i,L}} \log(n) \right] \right]}{r} \right). \end{aligned}$$

Notice that $\Gamma(\eta_i; 1) \geq 0$, $\log \left(\frac{\mu_i}{\phi_i \theta_i \kappa_{i,L} \lambda_{i,L}} \left(\frac{\theta_i \kappa_{i,U} DK}{r} \right)^{\frac{\lambda_{i,U}}{\lambda_{i,L}}} \vee 1 \right) \geq 0$. Therefore, we can write

$$\frac{c_i^*}{\lambda_{i,U}} \leq \log \left(\frac{\theta_i \kappa_{i,U} \left[\Gamma(\eta_i; 1) + \log \left(\frac{\mu_i}{\phi_i \theta_i \kappa_{i,L} \lambda_{i,L}} \left(\frac{\theta_i \kappa_{i,U} DK}{r} \right)^{\frac{\lambda_{i,U}}{\lambda_{i,L}}} \vee 1 \right) + 3 \frac{\lambda_{i,U}}{\lambda_{i,L}} \right] (n-1) \log(n)}{r} \right),$$

and after simplification we obtain the desired bound of:

$$c_i^* \leq \lambda_{i,U} \log \left(\frac{\theta_i \kappa_{i,U} C_U}{r} \right) + \lambda_{i,U} \log((n-1) \log(n)),$$

where $C_U = \Gamma(\eta_i; 1) + \log \left(\frac{\mu_i}{\phi_i \theta_i \kappa_{i,L} \lambda_{i,L}} \left(\frac{\theta_i \kappa_{i,U} DK}{r} \right)^{\frac{\lambda_{i,U}}{\lambda_{i,L}}} \vee 1 \right) + 3 \frac{\lambda_{i,U}}{\lambda_{i,L}}$. Observe that C_U does not depend explicitly on n , but through K it will be a function of parameters throughout the system.

Finally, it follows that $\lim_{n \rightarrow \infty} \frac{c_i^*}{\log(n)} \leq \lambda_{i,U}$.

- (ii) **Lower Bound:** We proceed with the lower bound identically. With our assumption of $f_i(x) \geq \kappa_{i,L} e^{-\frac{x}{\lambda_{i,L}}}$, we know

$$c_i^* \geq \lambda_{i,L} \log \left(\frac{\theta_i \kappa_{i,L} \left[\Gamma(\eta_i; 1) - (n-1) \log \left(\frac{\phi_i \theta_i}{\mu_i} \bar{F}_i(c_i^*) \right) \right]}{r} \right). \quad (\text{A.27})$$

Moreover, since $\Gamma(\eta_i; 1) \geq 0$ this term can be dropped to obtain:

$$c_i^* \geq \lambda_{i,L} \log \left(\frac{-\theta_i \kappa_{i,L} (n-1) \log \left(\frac{\phi_i \theta_i}{\mu_i} \bar{F}_i(c_i^*) \right)}{r} \right). \quad (\text{A.28})$$

By plugging in the initial crude bound of $c_i^* \hat{c}_i$, and since $\Gamma(\phi_i \hat{w}_i; 1) = -\log \left(\frac{\phi_i \theta_i}{\mu_i} \bar{F}_i(\hat{c}_i) \right)$ by definition, we can compute a tighter lower bound for c_i^* of

$$c_i^* \geq \lambda_{i,L} \log \left(\frac{\theta_i \kappa_{i,L} (n-1) \Gamma(\phi_i \hat{w}_i; 1)}{r} \right). \quad (\text{A.29})$$

This is precisely the lower bound in the first part of Lemma A.1. Note that for this result, we needed only the lower bound on $f_i(\cdot)$, through which (A.27) follows.

We now continue and prove the tighter lower bound, which requires the upper bound on $f_i(\cdot)$. In particular, we assumed that $f_i(x) \leq \kappa_{i,U} e^{-\frac{x}{\lambda_{i,U}}}$, and it follows that $\bar{F}_i(c_i^*) \leq \kappa_{i,U} \lambda_{i,U} e^{-\frac{c_i^*}{\lambda_{i,U}}}$. With (A.29), we can compute an improved upper bound of:

$$\bar{F}_i(c_i^*) \leq \kappa_{i,U} \lambda_{i,U} \left(\frac{r}{\theta_i \kappa_{i,L} (n-1) \Gamma(\phi_i \hat{w}_i)} \right)^{\frac{\lambda_{i,L}}{\lambda_{i,U}}}.$$

This upper bound on $f_i(\cdot)$ also implies that $\hat{c}_i \leq \lambda_{i,U} \log \left(\frac{\theta_i \kappa_{i,U} \Gamma(\eta_i; 1)}{r} \right)$. Similarly, the assumed $f_i(x) \geq \kappa_{i,L} e^{-\frac{x}{\lambda_{i,L}}}$ will give us $\bar{F}_i(\hat{c}_i) \geq \kappa_{i,L} \lambda_{i,L} e^{-\frac{\hat{c}_i}{\lambda_{i,L}}}$. Putting the two together, we will have

$$\bar{F}_i(\hat{c}_i) \geq \kappa_{i,L} \lambda_{i,L} \left(\frac{r}{\theta_i \kappa_{i,U} \Gamma(\eta_i; 1)} \right)^{\frac{\lambda_{i,U}}{\lambda_{i,L}}},$$

and it follows that

$$\bar{F}_i(c_i^*) \leq \bar{F}_i(\hat{c}_i) \frac{\kappa_{i,U} \lambda_{i,U}}{\kappa_{i,L} \lambda_{i,L}} \left(\frac{r}{\theta_i} \right)^{\frac{\lambda_{i,L}}{\lambda_{i,U}} - \frac{\lambda_{i,U}}{\lambda_{i,L}}} \frac{(\kappa_{i,U} \Gamma(\eta_i; 1))^{\frac{\lambda_{i,U}}{\lambda_{i,L}}}}{(\kappa_{i,L} \Gamma(\phi_i \hat{w}_i; 1))^{\frac{\lambda_{i,L}}{\lambda_{i,U}}}} (n-1)^{-\frac{\lambda_{i,L}}{\lambda_{i,U}}}.$$

Let $C_L = \frac{\kappa_{i,U} \lambda_{i,U}}{\kappa_{i,L} \lambda_{i,L}} \left(\frac{r}{\theta_i} \right)^{\frac{\lambda_{i,L}}{\lambda_{i,U}} - \frac{\lambda_{i,U}}{\lambda_{i,L}}} \frac{(\kappa_{i,U} \Gamma(\eta_i; 1))^{\frac{\lambda_{i,U}}{\lambda_{i,L}}}}{(\kappa_{i,L} \Gamma(\phi_i \hat{w}_i; 1))^{\frac{\lambda_{i,L}}{\lambda_{i,U}}}}$. Plugging this bound into (A.28), we obtain:

$$\begin{aligned} c_i^* &\geq \lambda_{i,L} \log \left(\frac{-\theta_i \kappa_{i,L} (n-1) \log \left(\frac{\phi_i \theta_i}{\mu_i} \bar{F}_i(\hat{c}_i) C_L (n-1)^{-\frac{\lambda_{i,L}}{\lambda_{i,U}}} \right)}{r} \right) \\ &\geq \lambda_{i,L} \log \left(\frac{-\theta_i \kappa_{i,L} (n-1) \log \left(C_L (n-1)^{-\frac{\lambda_{i,L}}{\lambda_{i,U}}} \right)}{r} \right), \end{aligned}$$

since $-\log \left(\frac{\phi_i \theta_i}{\mu_i} \bar{F}_i(\hat{c}_i) \right) = \Gamma(\phi_i \hat{w}_i; 1) \geq 0$, and hence this term can be dropped. Simplifying, we arrive at the desired bound of:

$$c_i^* \geq \lambda_{i,L} \log \left(\frac{\theta_i \kappa_{i,L} \lambda_{i,L}}{r \lambda_{i,U}} \right) + \lambda_{i,L} \log \left((n-1) \left[\log(n-1) - \frac{\lambda_{i,U}}{\lambda_{i,L}} \log(C_L) \right] \right),$$

from which it follows that $\lim_{n \rightarrow \infty} \frac{c_i^*}{\log(n)} \geq \lambda_{i,L}$.

Putting both (i) and (ii) together, we see that $c_i^* = \Theta(\log(n))$. \square

Proof of Proposition 4.3. Using Propositions 3.2 and 3.5, we can compute

$$\begin{aligned} \frac{V - \sum_{i=1}^n V_i}{T - t} &= \sum_{i=1}^n \left[-r(c_i^* - \hat{c}_i) + \theta_i (\bar{F}_i(\hat{c}_i) - \bar{F}_i(c_i^*)) [(n-1) + \Gamma(\eta_i; 1)] \right. \\ &\quad \left. + \theta_i (n-1) [\bar{F}_i(\hat{c}_i) \Gamma(\phi_i \hat{w}_i; 1) - \bar{F}_i(c_i^*) \Gamma(\phi_i w_i^*; 1)] \right], \end{aligned}$$

where V and V_i are evaluated at (t, x_1, \dots, x_n) and the difference becomes independent of wealths because of logarithmic utility. Notice that any of the terms in the sum will equal zero if $w_i^* = 0$ (in

which case we also must also have $\hat{w}_i = 0$, and hence $\hat{c}_i = c_i^*$. If not, then using the results from Section 4 we see that

$$\frac{-r(c_i^* - \hat{c}_i) + \theta_i (\bar{F}_i(\hat{c}_i) - \bar{F}_i(c_i^*)) [(n-1) + \Gamma(\eta_i; 1)]}{n} \xrightarrow{n \rightarrow \infty} \theta_i \bar{F}_i(\hat{c}_i),$$

since $c_i^* \asymp \log(n)$ and $\bar{F}_i(c_i^*) \rightarrow 0$. Moreover, we have seen that $(n-1)\bar{F}_i(c_i^*)\Gamma(\phi_i w_i^*; 1) = \Theta(1)$. Since the sum is now of order $|\mathcal{M}_n|$, putting the two together yields

$$\frac{V - \sum_{i=1}^n V_i}{T-t} = \Theta(n|\mathcal{M}_n|).$$

In Proposition 3.2, it is easy to see that $V_i = (T-t)\Theta(|\mathcal{M}_n|)$, and therefore we obtain

$$\frac{V}{\sum_{i=1}^n V_i} = 1 + \Theta(1),$$

as desired. \square

Proof of Corollary 4.4. This proposition is proved easily by analyzing the value functions in Propositions 3.2 and 3.5. We will use the notation of Section 4, where \hat{c} indicates the decentralized optimum, and c^* indicates the centralized optimum (likewise for w).

We begin by analyzing the decentralized value function V_i . Using the explicit formula in Corollary 3.3, we write:

$$\frac{V_i}{|\mathcal{M}_n|(T-t)} = \frac{J_i^*}{|\mathcal{M}_n|} + \frac{\log x}{|\mathcal{M}_n|(T-t)},$$

and see that the second term will go to zero as $n \rightarrow \infty$. Moreover, by assumption that all banks in \mathcal{M}_n are homogeneous, we will have $\hat{w}_{ij} = \hat{w}_{ik}$ for any $j, k \in \mathcal{M}_n$. This yields:

$$J_i^* = (1 - \hat{c}_i)r - \theta_i \bar{F}_i(\hat{c}_i)\Gamma(\eta_i; 1) + |\mathcal{M}_n| [\mu \hat{w} - \theta \bar{F}(\hat{c})\Gamma(\phi \hat{w}; 1)],$$

where \hat{c} denotes the optimal cash holdings for any bank in \mathcal{M}_n , and \hat{w} denotes the optimal investment made by any bank to those in \mathcal{M}_n . By using Eq (3.4) to compute \hat{w} , we obtain:

$$J_i^* = (1 - \hat{c}_i)r - \theta_i \bar{F}_i(\hat{c}_i)\Gamma(\eta_i; 1) + |\mathcal{M}_n| \left[\frac{\mu}{\phi} \left(1 - \frac{\phi \theta \bar{F}(\hat{c})}{\mu} \right) + \theta \bar{F}(\hat{c}) \log \left(\frac{\phi \theta \bar{F}(\hat{c})}{\mu} \right) \right],$$

and the desired limit follows. We note that this expression for J_i^* is only correct for $i \notin \mathcal{M}_n$, otherwise we would have a factor of $|\mathcal{M}_n| - 1$ in front of the term in brackets. However, in the limit this difference disappears.

The analysis of the centralized setting is almost identical, using the value function in Proposition 3.5, we have:

$$\frac{V}{n|\mathcal{M}_n|(T-t)} = \frac{J_C^*}{n|\mathcal{M}_n|} + \frac{\sum_{i=1}^n \log x_i}{n|\mathcal{M}_n|(T-t)}.$$

The only term of interest for large n will be J_C^* , and by homogeneity within \mathcal{M}_n we can see that:

$$J_C^* = |\mathcal{M}_n|(n-1)w^*\mu + \sum_{i=1}^n \left((1 - c_i^*)r - \theta_i \bar{F}_i(c_i^*) \left[\Gamma(\eta_i; 1) + (n-1)\Gamma(\phi_i w_i^*; 1) \right] \right),$$

where w^* denotes the optimal fractional amount invested into each bank in \mathcal{M}_n . Notice that only for bank in \mathcal{M}_n will we have c_i^* growing with n (logarithmically). Moreover, from the analysis in Section 4, we also know that $(n-1)\bar{F}_i(c_i^*)\Gamma(\phi_i w_i^*; 1)$ is of constant order. Therefore, when dividing by $n|\mathcal{M}_n|$ and taking the limit, the sum will go to zero. Only the first term will remain, and we also know that $w^* \rightarrow \phi^{-1}$ as $n \rightarrow \infty$, which concludes.

In order to show the limit for the price of anarchy, it is only necessary to sum V_i over n and divide. \square

B Price of Anarchy: Super-/Sub-Power Distribution

In this section, we perform similar calculations to the main result of Section 4, but for shock size densities bounded by power law distributions. In particular, we have the following analogue of Lemma A.1:

Proposition B.1. *If for all x we have $f_i(x) \geq \kappa_{i,L}(\zeta_i^0 + x)^{-\frac{1}{\alpha_{i,L}}}$, for some constants $\alpha_{i,L} < 1$, $\kappa_{i,L} > 0$, and $\zeta_i^0 \geq 1$, then*

$$c_i^* \geq \left(\frac{-\kappa_{i,L}\theta_i(n-1)\log\left(\frac{\phi_i\theta_i}{\mu_i}\bar{F}_i(\hat{c}_i)\right)}{r} \right)^{\alpha_{i,L}} - \zeta_i^0.$$

If, furthermore, the density satisfies $f_i(x) \leq \kappa_{i,U}(\zeta_i^0 + x)^{-\frac{1}{\alpha_{i,U}}}$, with $\kappa_{i,U} \geq \kappa_{i,L}$ and $\alpha_{i,L} \leq \alpha_{i,U} < 1$, then:

(i) **Upper Bound:**

$$c_i^* \leq C_U [(n-1)\log(n)]^{\alpha_{i,U}} - \zeta_i^0,$$

where C_U depends on all model parameters, but does not explicitly grow with n . As a result, $\lim_{n \rightarrow \infty} \frac{c_i^*}{[(n-1)\log(n)]^{\alpha_{i,U}}} \leq C_U$.

(ii) **Lower Bound:**

$$c_i^* \geq \left(\frac{\kappa_{i,L}\theta_i}{r}(n-1) \left[\left(\frac{\alpha_{i,L}}{\alpha_{i,U}} - \alpha_{i,L} \right) \log(n-1) - \log(C_L) \right] \right)^{\alpha_{i,L}} - \zeta_i^0,$$

for $C_L > 0$ depending only on i . Hence, $\lim_{n \rightarrow \infty} \frac{c_i^*}{[(n-1)\log(n)]^{\alpha_{i,L}}} \geq \left(\frac{\kappa_{i,L}\theta_i}{r} \left(\frac{\alpha_{i,L}}{\alpha_{i,U}} - \alpha_{i,L} \right) \right)^{\alpha_{i,L}}$.

The proof follows an identical technique. In the special case where the shock density is indeed a power distribution, we have the following analogue of Proposition 4.1.

Corollary B.2. *If $f_i(x) = \frac{(\frac{1}{\alpha_i}-1)(\zeta_i^0)^{\frac{1}{\alpha_i}-1}}{(\zeta_i^0+x)^{\frac{1}{\alpha_i}}}$, then*

$$c_i^* = \Theta\left([\log(n)]^{\alpha_i}\right).$$

This result can be seen by simply plugging $\alpha_{i,L} = \alpha_{i,U} = \alpha_i$ into Proposition B.1.

This Corollary can be used to replicate the remaining analysis in Section 4, but as the results are qualitatively similar, we omit these calculations.

B.1 Proof of Proposition B.1

Proof. The proof of this result largely mirrors the proof of Lemma A.1. Recall that we have shown that

$$\hat{c}_i \leq c_i^* \leq Kn^2,$$

for a suitable choice of K . By our assumptions on the density, it also follows that:

$$\left(\frac{y}{\kappa_{i,L}}\right)^{-\alpha_{i,L}} - \zeta_i^0 \leq f_i^{-1}(y) \leq \left(\frac{y}{\kappa_{i,U}}\right)^{-\alpha_{i,U}} - \zeta_i^0,$$

and

$$\frac{\kappa_{i,L}}{\frac{1}{\alpha_{i,L}} - 1} (\zeta_i^0 + x)^{1 - \frac{1}{\alpha_{i,L}}} \leq 1 - F_i(x) \leq \frac{\kappa_{i,U}}{\frac{1}{\alpha_{i,U}} - 1} (\zeta_i^0 + x)^{1 - \frac{1}{\alpha_{i,U}}}.$$

We can then follow the proof of Lemma A.1 identically, but using these bounds instead. \square

C Model Extensions

C.1 Liquidity Loss on Shock Arrival

In the main text of the paper, we make the assumption that the magnitude of a liquidity shock is important only if it exceeds the firm's cash holdings. We present an extension to our model that incorporates the magnitude of this loss directly into the dynamics of each firm's net capitalization.

Specifically, consider the event $d\tilde{N}_t^i = 1$, where a shock arrives for firm i . In the main text, this event induces a jump in firm i 's capitalization process X_t^i if and only if the size of the shock ζ_t^i is larger than the firm's available cash holdings c_t^i . However, we can build a model in which firms must themselves pay out any liquidity shocks, but only up to their maximum available supply of c_t^i . This feature is captured in the following equation for the dynamics of X_t^i .

$$\begin{aligned} \frac{dX_t^i}{X_t^i} = & \left((1 - c_t^i)r + \sum_{j \neq i} w_t^{ij} \mu_j + \frac{\eta_i}{\phi_i} \mu_i \right) dt - \sum_{j \neq i} w_t^{ij} \phi_j dN_t^j \\ & - (\eta_i + c_t^i) dN_t^i - \zeta_t^i (d\tilde{N}_t^i - dN_t^i), \quad i = 1, \dots, n. \end{aligned}$$

Notice that: i) when $dN_t^i = 1$, and a shock is larger than the firm's cash holdings their entire cash holdings are lost to cover the shock, and ii) if $d\tilde{N}_t^i = 1$ but $dN_t^i = 0$, and the shock was smaller than the firm's cash holdings, then the firm only loses exactly ζ_t^i – corresponding to their obligation to pay for the shock.

The value function for individual firm i , denoted $V_i(t, x)$ in the decentralized case, can be written identically to (3.1), but under the same regularity conditions of Proposition 3.1, the associated non-local PDE it solves is now given by:

$$\begin{aligned} 0 = \partial_t V_i + \sup_{c_i, w_i} & \left\{ \left[(1 - c_i)r + \sum_{j \neq i} w_{ij} \mu_j + \frac{\eta_i \mu_i}{\phi_i} \right] x \partial_x V_i + \theta_i \bar{F}_i(c_i) \left[V_i(t, x(1 - \eta_i - c_i)) - V_i \right] \right. \\ & + \sum_{j \neq i} \theta_j \bar{F}_j(c_j) \left[V_i(t, x(1 - \phi_j w_{ij})) - V_i \right] \\ & \left. + \theta_i \int_0^{c_i} f_i(u) \left[V_i(t, x(1 - u)) - V_i \right] du \right\}, \end{aligned} \tag{C.1}$$

with terminal condition $V_i(T, x) = U_i(x)$.

For simplicity, we will only consider the setting of a logarithmic utility function, (although the next step of separability can still be performed with power utility). We look for a separable solution to the PDE (C.1) given by $V_i(t, x) = g_i(t) + \log(x)$. By plugging this ansatz into (C.1) and simplifying, we can obtain the following ordinary differential equation for $g_i(t)$.

$$0 = g'_i(t) + \sup_{c_i, w_i} \left\{ (1 - c_i) r + \sum_{j \neq i} w_{ij} \mu_j + \frac{\eta_i \mu_i}{\phi_i} + \theta_i \bar{F}_i(c_i) \log(1 - \eta_i - c_i) \right. \\ \left. + \sum_{j \neq i} \theta_j \bar{F}_j(c_j) \log(1 - \phi_j w_{ij}) + \theta_i \int_0^{c_i} f_i(u) \log(1 - u) du \right\},$$

with terminal condition $g_i(T) = 0$.

Indeed the PDE solved by the value function has a separable solution, but now the first-order conditions for optimality solved by the optimal controls yield the following system of equations:

$$0 = \mu_j + \frac{\phi_j \theta_j \bar{F}_j(\hat{c}_j)}{1 - \phi_j \hat{w}_{ij}}, \quad \text{for all } j \neq i \\ 0 = -r - \theta_i f_i(\hat{c}_i) \log(1 - \eta_i - \hat{c}_i) - \frac{\theta_i \bar{F}_i(\hat{c}_i)}{1 - \eta_i - \hat{c}_i} + \theta_i f_i(\hat{c}_i) \log(1 - \hat{c}_i).$$

Although the second equation for the optimal control \hat{c}_i still exists in isolation, this expression cannot be solved explicitly for the optimal controls – certainly not for general choice of $\bar{F}(\cdot)$, and not even if we assume the liquidity shocks to be exponentially distributed. Hence, to study the optimal controls in this model, we would need to resort to computational methods.

Using the parameters in Table 4.1, and varying the value of λ , we can see in Figure C.1 that for large λ , and hence small expected size for liquidity shocks, the optimal cash holdings are largely identical. Since the first-order condition for the interbank investment amounts will only depend on the optimal cash holdings, this implies that the optimal allocations will coincide. As a result, the same results we obtain in the main text would hold.

C.2 Partially Liquid Risk-Free Asset

A further simplifying assumption made in this paper is that the risk-free asset, with constant rate of return r , does not provide a source of liquidity. In this section, we relax this assumption by allowing a firm's capital held in the risk-free asset to provide *some* liquidity upon shock arrival.

The only substantial difference we must make, with respect to the model in the main text, concerns the construction of the thinned Poisson process N_t^i from the shock arrival process \tilde{N}_t^i . We now define dN_t^i as follows:

$$dN_t^i = 1 \iff d\tilde{N}_t^i = 1 \text{ and } \zeta_t^i > c_t^i + \alpha \left(1 - c_t^i - \sum_{j \neq i} w_t^{ij} \right)_+,$$

where the parameter $\alpha \in (0, 1)$ controls the fraction of value obtained by liquidating the risk-free asset immediately, and $(y)_+$ denotes the positive part of y . Note that some firms may borrow at the risk-free rate ($1 - c_t^i - \sum_{j \neq i} w_t^{ij} < 0$), but obtain no additional liquidity for doing so.

With dN_t^i defined in this manner, the equation (2.2) for the dynamics of X_t^i remains identical. We define the value function of firm i in the same manner as the main text, and it can be shown




Figure C.1: Comparison of the optimal cash holdings when liquidity shocks directly induce losses regardless of size.

that under regularity conditions, $V_i(t, x)$ solves:

$$\begin{aligned}
0 = \partial_t V_i + \sup_{c_i, w_i} \left\{ \left[(1 - c_i) r + \sum_{j \neq i} w_{ij} \mu_j + \frac{\eta_i \mu_i}{\phi_i} \right] x \partial_x V_i \right. \\
+ \theta_i \bar{F}_i \left(c_i + \alpha \left(1 - c_i - \sum_{j \neq i} w_{ij} \right)_+ \right) \left[V_i(t, x(1 - \eta_i)) - V_i \right] \\
\left. + \sum_{j \neq i} \theta_j \bar{F}_j \left(c_j + \alpha \left(1 - c_j - \sum_{k \neq j} w_{jk} \right)_+ \right) \left[V_i(t, x(1 - \phi_j w_{ij})) - V_i \right] \right\}.
\end{aligned}$$

Indeed, if we assume firm i has logarithmic utility, and plug in the ansatz of $V_i(t, x) = g_i(t) + \log(x)$, the PDE admits a separable solution and we obtain the following ODE for $g_i(t)$.

$$\begin{aligned}
0 = g_i'(t) + \sup_{c_i, w_i} \left\{ \left[(1 - c_i) r + \sum_{j \neq i} w_{ij} \mu_j + \frac{\eta_i \mu_i}{\phi_i} \right] + \theta_i \bar{F}_i \left(c_i + \alpha \left(1 - c_i - \sum_{j \neq i} w_{ij} \right)_+ \right) \log(1 - \eta_i) \right. \\
\left. + \sum_{j \neq i} \theta_j \bar{F}_j \left(c_j + \alpha \left(1 - c_j - \sum_{k \neq j} w_{jk} \right)_+ \right) \log(1 - \phi_j w_{ij}) \right\},
\end{aligned} \tag{C.2}$$

with terminal condition $g_i(T) = 0$.

We note that the first order conditions may not be sufficient for the optimization problem in Eq. (C.2) since the objective function is not differentiable whenever $c_i + \sum_{j \neq i} w_{ij} = 1$. However, we can still differentiate the optimization problem with respect to the decision variables, and obtain

the following system of equations for the optimal controls of firm i :

$$\begin{aligned}
0 = & \mu_j - \theta_i f_i \left(\hat{c}_i + \alpha \left(1 - \hat{c}_i - \sum_{k \neq i} \hat{w}_{ik} \right)_+ \right) \left[-\alpha \cdot \mathbf{1} \left\{ 1 - \hat{c}_i - \sum_{k \neq i} \hat{w}_{ik} > 0 \right\} \right] \\
& - \frac{\phi_j \theta_j \bar{F}_j \left(\hat{c}_j + \alpha \left(1 - \hat{c}_j - \sum_{k \neq j} \hat{w}_{jk} \right)_+ \right)}{1 - \phi_j \hat{w}_{ij}} \quad \text{for all } j \neq i \\
0 = & -r - \theta_i f_i \left(\hat{c}_i + \alpha \left(1 - \hat{c}_i - \sum_{k \neq i} \hat{w}_{ik} \right)_+ \right) \log(1 - \eta_i) \left(1 - \alpha \cdot \mathbf{1} \left\{ 1 - \hat{c}_i - \sum_{k \neq i} \hat{w}_{ik} > 0 \right\} \right).
\end{aligned} \tag{C.3}$$

Observe that this system of equations differs from the system in Eq. (3.4) if and only if $1 - \hat{c}_i - \sum_{k \neq i} \hat{w}_{ik} > 0$, and hence if firm i holds a long position in the risk-free bond.

We have numerically obtained solutions for the optimal controls (\hat{c}_i, \hat{w}_i) in the symmetric case, where all banks have identical projects and utility functions. In doing so, we have seen that for large enough n , these optimal controls will exactly equal those in the main text (Section 3.2). The rationale for this is intuitive. Under the assumption of identical banks, the total amount invested in other banks' projects $\sum_{k \neq i} \hat{w}_{ik}$ grows linearly in n . This quantity eventually becomes large enough for bank i to short the risk-free bond and use this to finance their interbank project investments, which implies that the system of equations in Eq. (C.3) will simplify to the system of equations that gives rise to the optimal controls in Eq. (3.4).

We note that this observation will easily extend to the centralized optimum, since the planner will always choose to hold no less cash than decentralized banks. Hence under the planner's optimum, the optimal project investment amounts are no smaller than those in the decentralized setting, and it will still be the case that banks hold short positions in the risk-free bond.

Regardless, we have obtained numerical solutions for the optimal controls for decentralized banks. Figure C.2 compares the optimal cash holdings for a bank as α , the fraction of liquidity available due to a long position in the risk-free asset, varies. The parameters used in these simulations can be found in Table C.1. We note that even for a 2-bank system, cash holdings quickly drop to zero as α grows. And of course, for $\alpha = 0$, we obtain the optimal controls from the model in the main text. In Figure C.2b, with $n = 6$ we find that for the majority of values of α , the optimal cash allocations also coincide, since there is a sufficiently large number of return-bearing projects for a bank to invest in to justify shorting the risk-free asset.

Table C.1: Parameters used for obtaining optimal cash holdings shown in Figures C.2a and C.2b. Code is available [here](#).

Notation	Value	Description
r	0.03	Risk-free rate
μ	0.045	Excess drift
ϕ	0.4	Losses to External Investors
η	0.3	Losses to Associated Bank
$F(x)$	$1 - e^{-\frac{x}{\lambda}}$	CDF of shock size
λ	1	Parameter of $F(\cdot)$
θ	0.1	Shock arrival rate
γ	1	Relative risk aversion coefficient

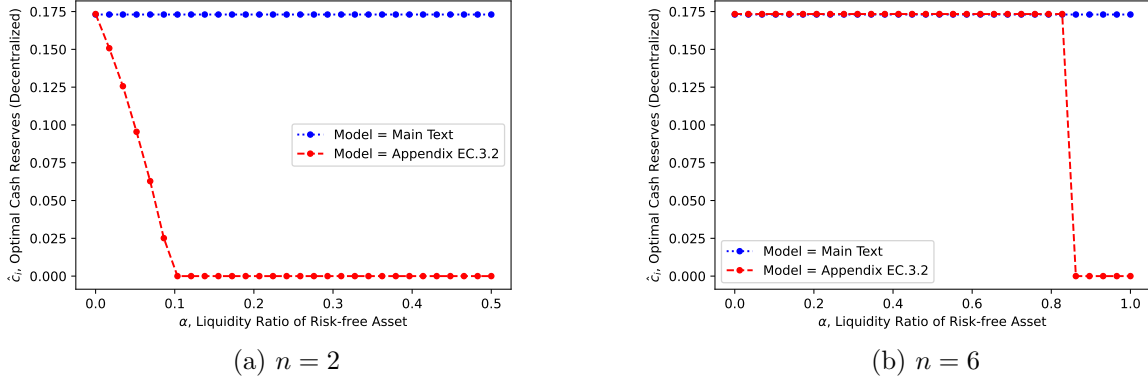


Figure C.2: Comparison of the optimal cash holdings when liquidity shocks directly induce losses regardless of size. Parameters (excluding n) are detailed in Table C.1.

C.3 Endogenizing Self-Investment Parameter η_i

A third assumption made in the main text concerns the fact that the parameter η_i , which governs the fraction of bank i 's wealth lost upon a liquidity shortage, is exogenous. In this section, we will relax this assumption and study the optimal control problem for bank i in a decentralized setting, when η_i can also be chosen freely.

The derivation of the HJB is largely identical, and after plugging in a separable form for the value function, we would obtain the following ODE for $g_i(t)$:

$$0 = g_i'(t) + \sup_{c_i, \eta_i, w_{ij}} \left\{ (1 - c_i)r + \sum_{j \neq i} w_{ij} \mu_j + \frac{\eta_i \mu_i}{\phi_i} + \theta_i \bar{F}_i(c_i) \log(1 - \eta_i) \right. \\ \left. + \sum_{j \neq i} \theta_j \bar{F}_j(c_j) \log(1 - \phi_j w_{ij}) \right\}. \quad (\text{C.4})$$

The only difference between this equation and its analogue in the main text is that η_i is included as a control variable. We will only analyze the optimization problem for both c_i and η_i , since the first-order conditions for each w_{ij} remain identical to that of the main text.

Now, the first-order conditions for these two controls are as follows:

$$0 = -r - \theta_i f_i(\hat{c}_i) \log(1 - \hat{\eta}_i), \quad 0 = \frac{\mu_i}{\phi_i} - \frac{\theta_i \bar{F}_i(\hat{c}_i)}{1 - \hat{\eta}_i}, \quad (\text{C.5})$$

where $\hat{\eta}_i$ denotes the optimal choice of η_i for bank i in the decentralized setting. Substituting $1 - \hat{\eta}_i$ by solving the second expression yields the following expression in only \hat{c}_i :

$$0 = -r - \theta_i f_i(\hat{c}_i) \log \left(\frac{\phi_i \theta_i \bar{F}_i(\hat{c}_i)}{\mu_i} \right).$$

In general, this expression cannot be explicitly solved without further assumptions on the distribution of liquidity shocks.

If, however, we assume that the size of liquidity shocks follows an exponential distribution, with $\bar{F}_i(x) = e^{-\frac{x}{\lambda_i}}$ and $f_i(x) = \lambda_i^{-1} e^{-\frac{x}{\lambda_i}}$, we can simplify and isolate \hat{c}_i to obtain:

$$\hat{c}_i = -\lambda_i \log \left(\frac{-r \lambda_i}{\theta_i \mathbb{W} \left(\frac{-r \lambda_i \phi_i}{\mu_i} \right)} \right),$$

where $\mathbb{W}(x)$ denotes the Lambert-W function defined as the inverse function of xe^x , restricted to the range $[-1, \infty)$ and the domain $[-e^{-1}, \infty)$. We note that this expression holds only when there is an interior solution to the optimization problem, i.e. when the optimal \hat{c}_i and $\hat{\eta}_i$ are strictly positive.

The analogue of this expression in the main text is:

$$\hat{c}_i = -\lambda_i \log \left(\frac{-r\lambda_i}{\theta_i \log(1 - \eta_i)} \right),$$

when the parameter η_i is assumed to be exogenous and given. The two are equal if and only if $\log(1 - \eta_i) = \mathbb{W} \left(\frac{-r\lambda_i\phi_i}{\mu_i} \right)$, or equivalently when $(1 - \eta_i) \log(1 - \eta_i) = \frac{-r\lambda_i\phi_i}{\mu_i}$.

We can numerically compute the value of the objective function for various choices of c_i , where we use the second expression in Eq. (C.5) to compute the optimal value of η_i for each possible c_i as follows:

$$\hat{\eta}_i(c_i) = \left(1 - \frac{\phi_i \theta_i \bar{F}_i(c_i)}{\mu_i} \right)_+,$$

since we assume that firm i still cannot short their own project. Figure C.3 shows, for different values of c_i , the quantity:

$$\text{Obj}(c_i) = (1 - c_i)r + \frac{\hat{\eta}_i(c_i)\mu_i}{\phi_i} + \theta_i \bar{F}_i(c_i) \log(1 - \hat{\eta}_i(c_i)), \quad (\text{C.6})$$

which corresponds to the portion of the objective function in Eq. (C.4) tied to the two controls η_i and c_i .

We have chosen the parameters in Table C.2 to show the existence of two optima. In one ‘corner’ solution, firm i chooses to hold zero cash and zero stake in their own project. In the second ‘interior’ solution, both of these optimal quantities are strictly positive fractions of i ’s wealth. This can be contrasted with the second curve, which shows Eq (C.6) when η_i is fixed at a value of $\frac{1}{2}$, and for which a unique maximum exists.

Note that the interior solution in Figure C.3 with endogenous η_i represents an equilibrium in which projects are *less* risky than their analogue in the main text model. Cash holdings are higher, and hence bank i ’s degree of investment in their own project is also greater than the fixed value of $\frac{1}{2}$. In contrast, the corner solution with zero cash holdings and zero investment in project i is not particularly interesting nor realistic, as it represents a setting in which bank i is wholly immune to all liquidity shocks.




Figure C.3: Comparison of the optimal cash holdings when liquidity shocks directly induce losses regardless of size.

Table C.2: Parameters used for generating Figure C.3. Code is available here.

Notation	Value	Description
r	0.013	Risk-free rate
μ_i	0.045	Excess drift
ϕ_i	0.8	Losses to External Investors
$F_i(x)$	$1 - e^{-\frac{x}{\lambda_i}}$	CDF of shock size
λ_i	1	Parameter of $F(\cdot)$
θ_i	0.1	Shock arrival rate
γ_i	1	Relative risk aversion coefficient