Formation of Optimal Interbank Lending Networks under Liquidity Shocks

Daniel E. Rigobon*
ORFE, Princeton

Ronnie Sircar†
ORFE, Princeton

November 17, 2022

Abstract

We formulate a model of the banking system in which banks control both their supply of liquidity, through cash holdings, and their exposures to risky interbank loans. The value of interbank loans jumps when banks suffer liquidity shortages, which can be caused by the arrival of large enough liquidity shocks. In two distinct settings, we compute the unique optimal allocations of capital. In the first, banks seek only to maximize their own utility – in a decentralized manner. Second, a central planner aims to maximize the sum of all banks' utilities. Both of the resulting financial networks exhibit a 'core-periphery' structure. However, the optimal allocations differ – decentralized banks are more susceptible to liquidity shortages, while the planner ensures that banks with more debt hold greater liquidity. 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 capitalization requirements that cause the decentralized system to achieve the planner's level of risk. In doing so, we find that systemically important banks must face the greatest losses when they suffer liquidity crises – ensuring that they are incentivized to avoid such crises.

*`drigobon@princeton.edu`
†`sircar@princeton.edu`

Contents

1 Introduction

Since the global financial crisis of 2008, systemic risk has become a topic of great interest to researchers, industry professionals, and policymakers alike. It is believed that interconnections between large financial institutions may have allowed distress to propagate throughout the financial system, and even beyond into other economic sectors. The crisis was extremely costly – even with bailouts of nearly \$500 billion (provided by the US government), losses to the global economy totaled over \$2 trillion (Lucas, 2019). This event renewed researchers' interest in understanding fragility of financial systems, and how policymakers can effectively intervene.

The phenomenon in which distress spreads through an entire system is dubbed 'systemic risk'. In financial applications, it is natural to study systemic risk through the perspective of networks and complex systems – the spread of distress is facilitated by the financial network that links institutions. Its edges represent, for instance, interbank loans (Allen and Gale, 2000) or overlapping portfolios (Cifuentes et al., 2005), and thereby highlight pathways along which distress can propagate. We may therefore expect that the observed patterns of financial contagion are related to this underlying network structure.

There is a substantial amount of research studying the dependence of systemic risk on characteristics of the financial network; for recent surveys see Jackson and Pernoud (2021) and Benoit et al. (2017). For example, a well-known contribution by Allen and Gale (2000) studies how several stylized networks of direct interbank claims can yield different patterns of contagion – or even no contagion at all. The initial shock caused by idiosyncratic demand for liquidity can cause an overwhelmed banks' neighbors to suffer further liquidity shortages. A fully connected network is found to be optimal for sharing liquidity and therefore reducing the possibility of a systemic crisis. However, Gai and Kapadia (2010) identify a substantial tradeoff – the resilience of highly connected financial networks is accompanied by an increased intensity of systemic events. This feature, dubbed 'robust yet fragile', implies that although systemic crises are unlikely, they cause catastrophic and widespread damage.

A critical assumption in this branch of the literature is that the financial network is *exogenous*. Such papers are therefore restricted to analyzing the effect of a particular generative model for the network on systemic risk. Although these models serve as a useful baseline, this assumption is unlikely to hold in practice. Instead, the network structure may be *endogenous*; each financial institution's connections reflect a set of optimal decisions. This perspective has become more prevalent in the literature, with relevant contributions by Bluhm et al. (2014), Acemoglu et al. (2015), and Farboodi (2021). Interestingly, it is possible for systemic risk to emerge hand-in-hand with each financial institution's selfish optimal behavior. However, these individually-optimal decisions need not maximize the collective well-being of the financial system. In such cases, as in this paper, financial institutions may be failing to internalize the negative effects of their decisions on the entire system. It is therefore of interest to analyze the severity of these negative externalities and how they might be remedied.

In this paper, we study the formation of such a continuous-time interbank lending network when banks both face (and can insure against) idiosyncratic liquidity shocks. We formulate a system-wide optimization problem for both interbank exposures and insurance, in which the resulting network of interbank linkages indicates the channels for (and magnitude of) the propagation of financial distress.

The model proceeds as follows: consider a financial system comprised of a given number of banks. These 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 (i.e. a portfolio of commercial loans) available to each bank. We

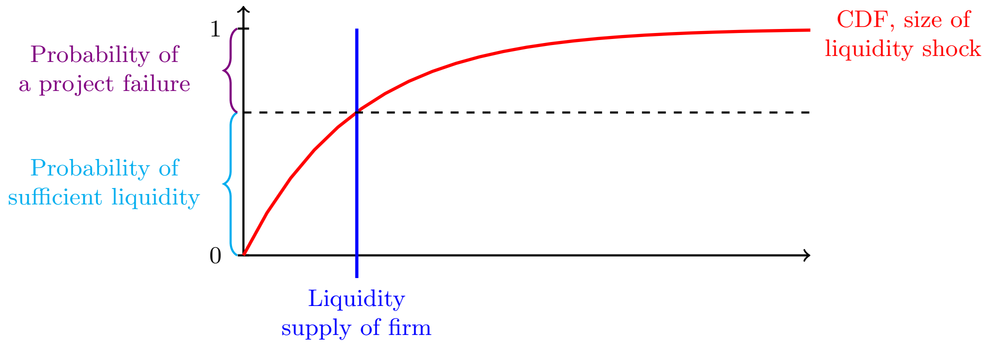

Figure 1.1: For a single bank, the relationship between the cumulative distribution function (CDF) of the size of a liquidity shock, their supply of cash, and the conditional probabilities of (in)sufficient liquidity.

assume that these opportunities are scalable, but are only accessible to their associated bank. The interbank lending mechanism in our model allows, for example, a deposit-collecting bank A to obtain the large returns of investment bank B's unique opportunities through a direct loan of capital from A to B – after which B invests this amount into their revenue-generating operations. In this setup, bank B is effectively operating as an intermediary between bank A and B's own investment opportunities. We note that this construction is similar to the model of both Rochet and Tirole (1996) and Acemoglu et al. (2015), wherein banks invest in each other's 'projects' (henceforth, we will also use 'projects' to refer to these unique investment opportunities). In both these models and ours, the riskiness of these projects is tied to some decision 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 is periodically struck by liquidity shocks of random magnitude, and if the size of a shock exceeds the bank's current supply of liquidity (i.e. cash), then the project's value instantaneously drops (we refer to this event as a 'project's failure'). These shocks are assumed to represent, for example, additional liquidity required for the project to succeed, such as occurs in Rochet and Tirole (1996), Acemoglu et al. (2015), and more. If the required amount of cash is available, then the project continues smoothly. Conversely, a project's failure results in all its investors suffering losses proportional to their stake. Therefore, conditioned on the arrival of a liquidity shock, a bank's supply of cash determines their project's level of risk. Figure 1.1 illustrates the relationship between a bank's liquidity supply, the distribution of a liquidity shock's size, and the probabilities of each outcome.

In the model, bank A is assumed to have non-zero stake in their own project, and is therefore a co-investor of its creditors. Without this assumption, bank A would have no incentive to hold liquidity – they would be unaffected by their project's failure. In addition to holding cash, recall that bank A may lend their capital to any other bank B, which is invested into B's project. Banks in the system may also invest in a risk-free bond, or borrow at this rate from the central bank or external financiers. Finally, each bank is assumed to have some fixed amount of deposits, which fully specifies their balance sheet. An example is given in Figure 1.2, with descriptions of each item.

A key focus of this paper is that each bank endogenously chooses to allocate its capital between cash (i.e. supply of liquidity) and risky interbank loans. To that end, we will study the optimal capital allocations for two extreme settings of the financial system. First, 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 some terminal time. We note that this setting reflects a game-

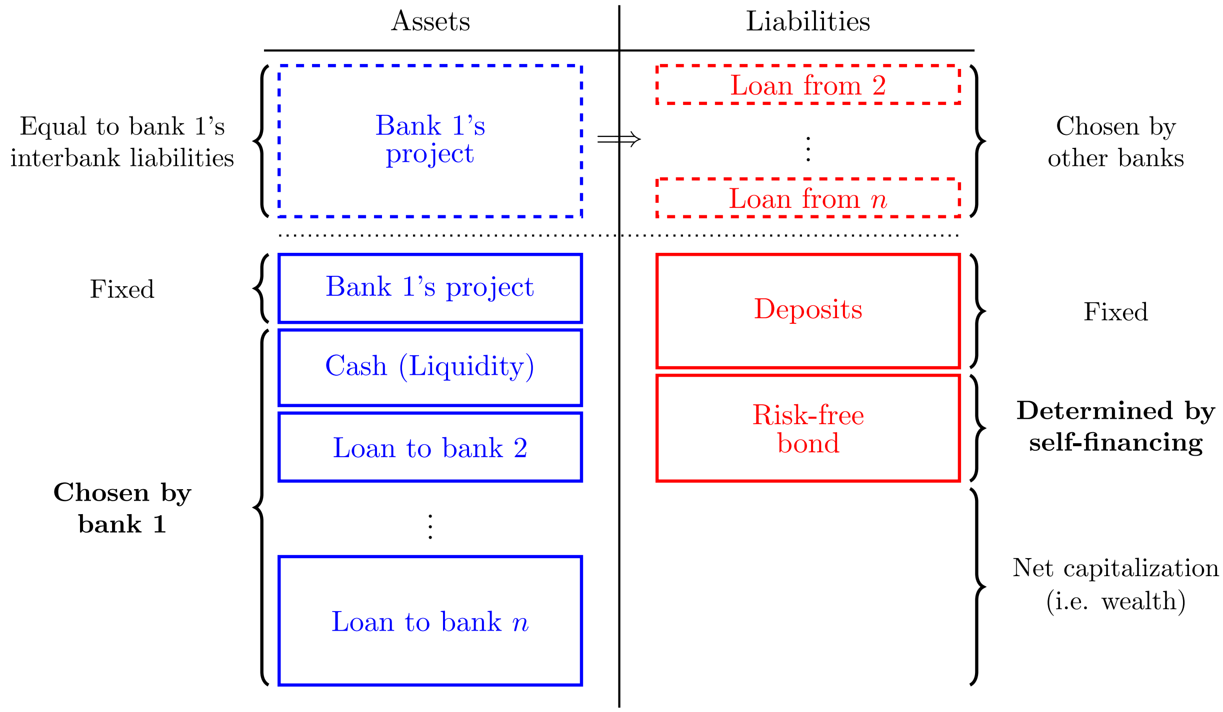

Figure 1.2: An example balance sheet for bank 1, who borrows at the risk-free rate to finance their investment portfolio. The decision variables for the bank are highlighted in bold.

theoretic equilibrium. Second, we consider the *centralized* setting – where a single central planner makes the allocation decisions for all banks concurrently. The planner aims to maximize the sum of individual banks' utilities. In both of these cases, we derive the dynamic programming equations for the respective value functions, and explicitly compute the optimal allocations. Under stricter conditions, we can conclude uniqueness.

We observe a discrepancy between the optimal allocations computed in both settings; the central planner often chooses to hold a greater supply of liquidity. This occurs because our model captures a simple negative externality; when individualistically choosing their cash holdings, a bank determines the risk experienced by its creditors. An individual bank, operating in a selfish manner, fails to consider its creditors' losses when choosing their supply of cash. In contrast, the planner is cognizant of this systemic effect and acts accordingly by reducing the level of risk for banks with larger debt. Namely, the central planner achieves the welfare-maximizing (i.e. first-best) allocation for the financial system. As a consequence of this discrepancy, a project's failure is more likely to occur with decentralized behavior than with a central planner. However, we also observe that the size of interbank loans is larger in the centralized setting, and hence each project's failure becomes more damaging to the system. This tension between the likelihood and severity of failures bears a resemblance to the 'robust yet fragile' observation made by Gai and Kapadia (2010). 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 has been dubbed the 'price of anarchy' by 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' co-investment requirements to replicate the planner's optimal allocation.

There are several interesting consequences of our paper. First, we find that the central planner's optimal allocation leads to low-frequency and high-intensity losses to the system. This may imply that the 'robust-yet-fragile' feature of financial networks is socially optimal. However, we see that the planner perfectly compensates for the larger-magnitude losses by ensuring they are less likely. As a result, the centralized allocation involves greater lending throughout the system. Additionally, in both settings we see that the (optimal) endogenous financial networks exhibit a strong 'core-periphery' structure, where only a subset of banks serve as borrowers to the rest of the system. Intuitively, we also show that systemically important banks must face the greatest losses if they are to replicate the planner's optimal allocation. This lends credence to the perverse incentives caused by 'too-big-to-fail' policies or other government bailouts.

This paper is organized as follows. Section 1.1 reviews several relevant branches of literature. Section 2 introduces the model of interbank lending and the dynamics of each financial instrument. In the first part of our main results, Section 3 derives the optimal allocation in the decentralized (Section 3.1) and centralized (Section 3.2) 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

The foundational papers on continuous-time portfolio optimization are by 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 techniques we use in this paper for deriving the optimal allocation will be similar to Merton's original work and its subsequent branch of literature. However, 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 to control the jump intensity of a risky asset's returns has not been previously studied in the area of portfolio optimization.

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 Possamaï (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 supply of cash 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 individuals optimize greedily. Additionally, our setting assumes perfect information.

A strong motivation for this paper follows from the systemic risk literature; 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. More recent 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 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. 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. 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 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 edges are cost-less to form.

Most closely related to this paper is the study of endogenous financial networks, including Zawadowski (2013); Bluhm et al. (2014); Acemoglu et al. (2015); Babus (2016) and Farboodi (2021). 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 see that individual banks' optimal behavior fails to internalize 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 selfish optimal decisions.

The three papers most similar to our own are Bluhm et al. (2014),Acemoglu et al. (2015) and Farboodi (2021). Our model is mechanistically different from the models in these studies – which are either static or consist of three distinct time periods. In contrast, we analyze the optimization problems in a dynamic continuous-time environment. First, Bluhm et al. (2014) construct a model of optimal interbank lending where banks face both liquidity and capital 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. Despite the similarities to our paper, the authors largely focus on numerical and simulation results. In contrast, we seek to provide a theoretical characterization of the optimal solution wherever possible. Moreover, our model endogenizes the initial sources of disruption.

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, an interbank loan will also allow the lending bank to access the surplus generated by a risky investment of the borrowing 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 many 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 can, however, 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. (2015) endogenize both the decision of interbank lending and also the interbank interest rates. In a similar spirit to Rochet and Tirole (1996), 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 model of interbank lending is 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 lending opportunities.

2 Model

Consider a financial system consisting of n different banks. Let $(\Omega, \mathcal{E}, \mathbb{P})$ be a probability space, containing n independent Poisson processes $\tilde{N}_t^1, ..., \tilde{N}_t^n$, $t \geq 0$, each of which has corresponding intensity $\theta_1, ..., \theta_n > 0$. 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 $\{X_t^i\}_{t\geq 0}$. We now aim to describe the dynamics of a bank's wealth.

2.1 Dynamics of Interbank Loans

First, the financial system contains a risk-free bond, which accumulates a constant, fixed rate of return r. Therefore its price, denoted S_t^0, evolves according to the ordinary differential equation $\frac{dS_t^0}{S_t^0} = rdt$. Banks may both borrow and invest at this risk-free rate, but in our model, we assume that an investment in the bond does not provide liquidity.

Each bank i has access to a unique set of external investments, e.g. a collection of commercial loans (henceforth referred to as a 'project'). These projects are available to another bank $j \neq i$ in the system through an interbank loan provided to i. In this manner, bank i serves as an intermediary between its project and the lending bank j. We will assume that there is no fee associated with this intermediation. Additionally, the capital invested in interbank loans is assumed to be illiquid. More precisely, neither the lending nor borrowing banks can use capital invested in a project to meet their liquidity needs. While these interbank claims indirectly accumulate large constant rates of return for investors, they will incur losses when the borrowing bank's project fails. If such a failure occurs, then the value of all capital invested in the project immediately drops. For instance, it is plausible that a bank's revenue operations intermittently require additional liquidity to cover a position or meet regulatory requirements. A failure to do so may lead to an inability to realize an investment's gains, or even directly cause losses.[1]

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

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

Observe that since $\mu_i > 0$, this interbank claim has rate of return larger than r. The jump increment dN_t^i is obtained by performing a thinning of the shock arrival process $\tilde{N}_t^i$, and is described in the next subsection. The increment takes on values in $\{0, 1\}$, and is non-zero if and only if bank i's project fails at time t. Finally, ϕ_i represents the magnitude of losses borne by investors in a project when it fails, i.e. $1 - \phi_i$ is the recovery rate.

2.1.1 Liquidity Shocks and Risk

All projects in the system may experience liquidity shocks; if sufficiently large, these shocks induce the project's failure. A key feature of this paper is each bank's ability to control their project's susceptibility to failure – by holding a greater supply of liquidity, banks' projects are safer. In our model, this is represented through bank i's ability to influence the intensity of the jump increment dN_t^i that appears in (2.1).

A bank may hold a non-negative amount of their capital as cash, which has a constant price of 1. Although this capital effectively depreciates at the risk-free rate r (as it cannot be invested in the bond), it is the *only* source of liquidity within the system, and is the sole manner in which a bank can hedge against liquidity shocks. Namely, if a liquidity shock exceeds bank i's supply of cash, their project experiences a failure, and investors incur losses. The jump increment dN_t^i in (2.1) represents the arrival of shocks that overwhelm bank i's supply of cash. Its construction follows from a probabilistic model of liquidity shocks and a bank's supply of cash.

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 $X_t^i \zeta_t^i$, 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) $F_i(\cdot)$. The complementary CDF of ζ_t^i is defined as $\bar{F}_i(\cdot) = 1 - F_i(\cdot)$.

Let $c_t^i \geq 0$ denote the fraction of bank i's capital held in cash at time t. When the shock to bank i is larger than their supply of liquidity (i.e. $\zeta_t^i > c_t^i$), their project fails. When this occurs, all investors in the project suffer an instantaneous return of $-\phi_i$ on their investment amount. In particular, if $c_t^i = 0$, then any liquidity shock to bank i at time t, no matter how small, results in their project's failure.

The jump process N_t^i is constructed by independently flipping coins at every arrival time of $\tilde{N}_t^i$, with success probability given by $p_t^i = \bar{F}_i(c_t^i)$. Observe that $p_t^i = \mathbb{P}\left(\zeta_t^i > c_t^i \,\middle|\, \tilde{N}_t^i = 1\right)$. If and only if the flip is won, we let $dN_t^i = 1$. It follows that the instantaneous rate (at time t) of the Poisson

[1] There are several other distinct justifications for this feature of our model. For example, the bank may be investing in costly, continuous monitoring of its project, which reduces the risk of it suffering losses (e.g. default of its commercial loans). A different interpretation considers the effect of consumers' random liquidity preference. We may imagine that the shock represents depositors' demand to withdraw cash, and if the bank fails to meet this demand, they must liquidate the project at a loss to meet their more senior obligation.

process N_t^i is equal to $\theta_i \bar{F}_i\left(c_t^i\right)$.[2] 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's project fails, conditional on the time-t arrival of a liquidity shock with CDF $F_i(\cdot)$. See Figure 1.1 for an illustration.

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 a large enough supply of liquidity. Since the return of this project would be greater than the risk-free rate, this would lead to all other banks to profit infinitely by borrowing at the risk-free rate and investing in the 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 financial network.

2.2 Dynamics of Wealth

In this model, a bank may provide interbank loans to another; let $w_t^{ji} \geq 0$ denote the fraction of bank j's capital lent to bank $i \neq j$. The return experienced by this interbank claim is given by (2.1). Recall that c_t^i equals the fraction of bank i's wealth held as cash, which accumulates no return over time. Therefore, the remaining $X_t^i(1 - c_t^i - \sum_{j \neq i} w_t^{ij})$ units of wealth are invested in (or borrowed at) the risk-free rate.

The final component influencing bank i's wealth is their degree of investment in their own project. We assume that each bank invests a fixed, given fraction of their current wealth. Unlike the interbank loans, we will assume that this quantity cannot be controlled.[3] This assumption has several possible interpretations. First, it may be the case that bank i is required by its creditors to be a co-investor in its project. We may also imagine that these projects are initialized by their respective banks, and simply scaled by any additional investments from the rest of the system. Therefore, the cost of initialization must be borne by the borrowing bank.

We will use $\frac{\eta_i}{\phi_i}$ to denote the fraction of 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 their project fails.[4] The parameter η_i captures the severity of a project's failure on the associated bank – in the extreme case of $\eta_i = 1$, a single failure will wipe out the bank i. Conversely, if $\eta_i = 0$, then bank i has no stake in their project and is unaffected by its failure. We will take $\eta_i \in (0, 1)$, away from the two extreme cases.

Putting together the dynamics for each component of bank i's wealth, we see that X_t^i, follows

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

[2] This 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.

[3] In principle, we could imagine allowing bank i to also control their exposure to their own project, while only being subjected to a minimum requirement. However, doing so introduces significant challenges in characterizing the optimal allocations.

[4] If, instead, this were a fixed amount and not fraction, then as a bank's wealth grows, their incentive to hold liquidity would become weaker. Such a setting is quite interesting in its own right, and may perhaps lead to a cyclic supply of liquidity – but it is not the focus of this paper.

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

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

A novel contribution of this paper is the control c_t^i – while there is no return accumulated by this capital held as cash, it serves to reduce the likelihood that bank i's project fails, which would cause them to lose a fraction η_i of their wealth.

We say that $(c_\cdot^i, w_\cdot^{i\cdot}) \in \mathcal{A}_{s,t}^i$, 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.

All banks seek to maximize their own utility of wealth at a common terminal time $T < \infty$. 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:

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

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 3.2 will consider the perspective of a single central 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 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 utility that could exist in the financial system if banks coordinated. 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, which represents direct balance sheet exposures between banks.

3.1 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:

$$V_i(t,x) = \sup_{(c_\cdot^i, w_\cdot^{i\cdot}) \in \mathcal{A}_{t,T}^i} \mathbb{E}\left[U_i(X_T^i) \middle| X_t^i = x \right]. \tag{3.1}$$

Recall that $A_{t,T}^i$ denotes the set of admissible controls for bank i – defined in Section 2.2. Note also that each bank is simultaneously solving their own optimization problem, and therefore the value function in (3.1) of bank i may depend on the allocations chosen by other banks within the system. In this sense, the value functions are related and our model's setup can be considered game-theoretic.

Our first result derives the non-local dynamic programming equation (which is often referred to as the Hamilton-Jacobi-Bellman 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):*

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

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. Assuming existence of the optimal controls is verified by 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:

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

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 wealth. More precisely, $\Gamma(\delta;\gamma_i) = x^{\gamma_i - 1}\left[U_i(x) - U_i(x(1-\delta))\right]$ for any $x > 0$.

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

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

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

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) otherwise, 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 also 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 $c_i, w_{i\cdot}$.

Remark 3.1. *The optimal interbank loan $\hat{w}_{ij}$ depends explicitly on $\hat{c}_j$ through the function $\bar{F}_j(\hat{c}_j)$. Moreover, for any choice of bank j's cash supply, 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 choice. However, bank j's optimal value $\hat{c}_j$ 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 bank' 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, this result verifies the assumption made in Proposition 3.1 regarding the existence of optimal controls.

Analysis of Decentralized Optimum

With explicit solutions for the optimal allocations, it is possible to analyze their dependence on the exogenous parameters of the system. First, note that the optimal interbank loan $\hat{w}_{ij}$ depends on bank 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 interbank loans are equal, the nominal amounts may differ. However, the optimal lending amount is decreasing in the lender's risk aversion coefficient γ_i, as we might expect.

From (3.4), we can also see that $\hat{c}_i$ is decreasing in the risk-free rate. This occurs because cash is effectively depreciating at the risk-free rate r. However, each unit of additional cash provides a marginal benefit by lowering the risk of a bank's project failing. From the proof of Proposition 3.2, the optimal choice of $\hat{c}_i$ will solve the following:

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

which indicates that the resulting $\hat{c}_i$ achieves the optimal tradeoff between the cost of liquidity and induced risk. In particular, the optimal $\hat{c}_i$ ensures that the marginal cost of holding liquidity (r)

equals the marginal benefit of reducing risk $(\theta_i f_i(\hat{c}_i)\Gamma(\eta_i;\gamma_i))$. 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, i.e. $\hat{c}_i = 0$.

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 returns for bank j's project can be controlled by bank j itself. Nonetheless, notice that the optimal investment $\hat{w}_{ij}$ grows with this 'Sharpe-like' ratio. If, in particular, the ratio is less than one, then the expected excess return of the interbank loan (equal to $\mu_j - \phi_i\theta_i\bar{F}_j(\hat{c}_j)$) is negative, and bank i would in fact 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 bank j has creditors 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 a subset of banks serve as the only borrowers – an example of such a financial network can be seen in Figure 3.1.

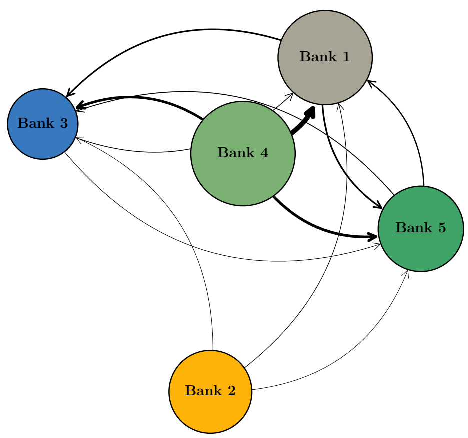

Figure 3.1: Sample financial network generated by the decentralized optimum. Edges point from lending to borrowing banks, and their width indicates the nominal size of the exposures. Node size indicates to total capitalization.

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	$e^{-2.4x}$

Table 3.1: Parameters for the financial network in Figure 3.1. The risk-free rate is equal to $r = 5\%$. Code generating this figure can be found here.

3.2 Centralized Network

Consider now the perspective of a single central planner of the financial system. In contrast with Section 3.1, we will see that the planner has two different incentives for bank i's holding of cash. The first is identical – bank i stands to lose wealth if their project fails. The second incentive is systemic – other banks face losses on their interbank claims upon the very same event. Therefore, we expect the planner to have stronger incentives to hold cash, and elect for a greater supply of liquidity within the system.

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

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

where $\mathcal{A}_{t,T} = \prod_i \mathcal{A}^i_{t,T}$ 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.*

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

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

This inequality reflects an inefficiency of the decentralized setting; the planner's allocation is the first-best (i.e. welfare-maximizing) outcome for the system. 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 $\mathcal{C}^{1,1,\dots,1}([0,T), \mathbb{R}_+, ..., \mathbb{R}_+)$, then it solves*

$$\begin{aligned} 0 = \partial_t V + \sup_{c_\cdot, w_{\cdot\cdot}} \Bigg\{ \sum_{i=1}^{n} \Bigg(&\left[(1-c_i)\, r + \sum_{j \neq i} w_{ij}\mu_j + \frac{\eta_i \mu_i}{\phi_i} \right] x_i \partial_{x_i} V \\ &+ \theta_i \bar{F}_i(c_i) \Big[V(t, x_1(1-\phi_i w_{1i}), .., x_i(1-\eta_i), .., x_n(1-\phi_i w_{ni})) - V \Big] \Bigg) \Bigg\}, \end{aligned} \tag{3.7}$$

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

The proof of this result is only a minor adaptation to the proof of Proposition 3.1, and can 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 project fails, 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 incurred by such a failure.

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, to establish uniqueness, we will need the following technical assumption.

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

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

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

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

Assumption 2 is sufficient for uniqueness of the planner's the optimal allocation. While we have numerically observed that the optimal solution is almost always unique, 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.

Analogous to Section 3.1, 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 and lending amounts for the planner, which solve the following system of equations:*

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

Letting c_i^* and $w_{\cdot i}^*$ be the optimal allocation, we define

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

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

$$V(t, x_1, \ldots, x_n) = g(t) + \sum_{i=1}^{n}\log x_i. \tag{3.11}$$

Furthermore, under Assumption 2, the optimal cash and lending amounts $(c_i^*, w_{\cdot i}^*)$ 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 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 $\{c_i, w_{1i}...w_{ni}\},\ i = 1...n$ can be analyzed separately, which greatly simplifies our analysis. However, the coupling between c_i and $w_{\cdot i}$ 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 $w_{\cdot i}$ is strictly concave and admits a unique solution (these can be seen in the proof of Proposition 3.5). Conversely, for given values of $w_{\cdot i}$, 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 unique optimal allocation.

Since we have shown existence of an optimal allocation, and the proposed solution in (3.11) 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* (3.11)*. Furthermore, the optimal interbank lending and cash amounts solve* (3.4)*.*

Analysis of Centralized Optimum

There is one main difference between the system of equations in (3.10) and the optimal solutions in (3.4) obtained from the decentralized setting. Here, we have an additional term of $(n-1)\Gamma(\phi_i w^*_{\cdot i}; 1)$ that influences the planner's optimal value for c^*_i. This term directly captures the externality – when i's project fails, the planner sees losses in utility experienced by all banks. As a result, with more banks the planner maintains a larger supply of cash to compensate for greater systemic 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 $w^*_{\cdot i} \geq 0$ in (3.10), the planner has a greater incentive to hold liquidity than the individual bank.[5] Hence, the planner will hold more liquidity than the decentralized optimal allocation – we will study this difference more closely in the following section. Finally, we also notice that given the amounts of cash held, the optimal investments $w^*_{\cdot\cdot}$ and $\hat{w}_{\cdot\cdot}$ are computed identically. It follows that any differences between the optimal interbank lending amounts in (3.4) and (3.10) can only be driven by differences in optimal cash supplies.

4 Price of Anarchy

It is natural to compare the two optimal allocations from Sections 3.1 and 3.2. In particular, we may be interested in computing the gap in welfare from the inequality (3.6). More generally, in simulations we see stark differences between the two optimal allocations. Figure 4.1 illustrates a sample path for the wealth of three banks, where in 4.1a the controls are given by (3.4), and in 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 $\hat{w}_{\cdot i}$ (this follows from $\gamma_j = 1$ for all j). Additionally, recall that $c^*_i, w^*_{\cdot i}$ 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 $A_1 \leq \lim_{n\to\infty} \frac{g(n)}{h(n)} \leq A_2$. If $A_1 = A_2$, then we will write $g(n) \asymp h(n)$.

4.1 Liquidity Supply and Project Risk

Comparing the two optimal allocations, since $w^*_{\cdot i} \geq 0$, it will necessarily be the case that $c^*_i \geq \hat{c}_i$. Our fundamental result establishes the asymptotic rate at which the planner's optimal supply of liquidity grows as the size of the system increases. More precisely, we show that for heavy-tailed distributions, the planner's supply of cash must grow at least logarithmically in the system size n

[5] This 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 fails. However, clearly this result may not align with the best outcome for bank i itself.

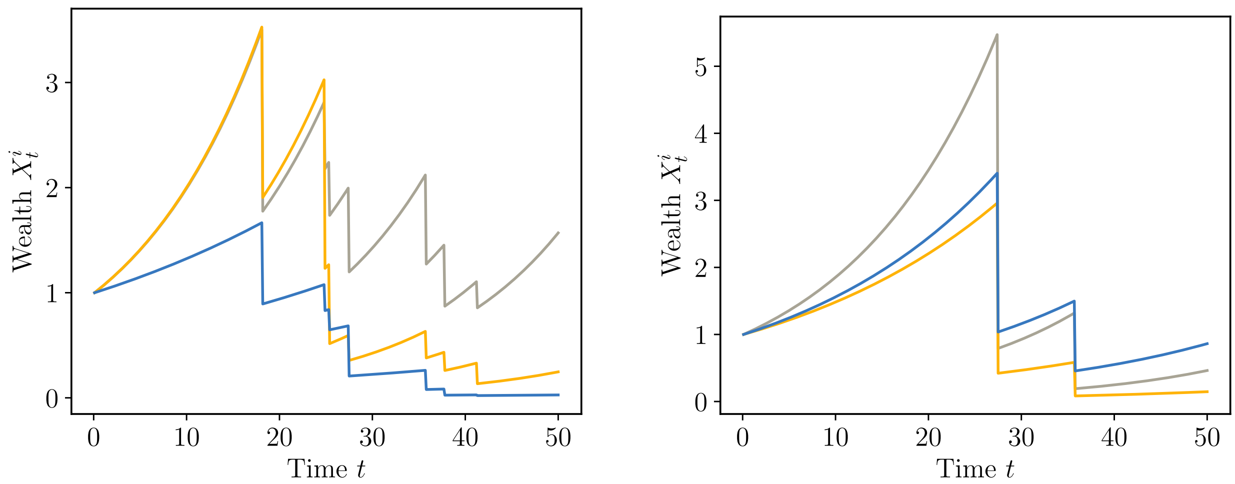

(a) Decentralized

(b) Centralized

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.

– and under stronger assumptions this lower bound is tight. In contrast, if $w_{\cdot i}^* = 0$, then we would have $c_i^* = \hat{c}_i$ – which is of constant order.

Proposition 4.1. *Assume that $w_{\cdot i}^* > 0$. 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*

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

In particular, the planner's optimal cash supply asymptotically grows 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:***

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

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, $\lim_{n\to\infty} \frac{c_i^*}{\log(n)} \leq \lambda_{i,U}$.

(ii) ***Lower Bound:***

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

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

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

The proof is provided in Appendix A.3. It follows from iterating through upper (and lower) bounds for c_i^* using the system of equations in (3.10), and beginning from crude estimates.[6]

The following special case of Proposition 4.1 occurs when the shock sizes are exponentially distributed.

Corollary 4.2. *If* $F_i(x) = 1 - e^{-\frac{x}{\lambda_i}}$*, then*

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

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

The proof follows from plugging in $\lambda_{i,L} = \lambda_{i,U} = \lambda_i$, and $\kappa_{i,L} = \kappa_{i,U} = \lambda_i^{-1}$. Simplifying the constant C_L that appears in the lower bound yields the desired result.

Corollary 4.2 is a useful tool for comparing the two optimal allocations, as all differences are driven by the distinct supply of liquidity. From here onward, we assume the setting of this Corollary, wherein all shock sizes are exponentially distributed. First, we can directly compute the dependence of project i's likelihood of failure on the system size. We see that:

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

and it follows that $\bar{F}_i(c_i^*) = \Theta\left(\frac{1}{n\log(n)}\right)$.[7] In stark contrast, the optimal intensity from the decentralized setting, $\bar{F}_i(\hat{c}_i)$, is constant in n. That is, $\bar{F}_i(\hat{c}_i) = \Theta(1)$. These two results will allow us to analyze the price of anarchy.

4.2 Losses to Lending Banks

In addition to the supply of liquidity, the optimal investment amounts will differ between the two settings. Due to the greater risk of jumps in the decentralized optimum, banks will invest less capital into each others' projects, and have a lesser degree of integration with the system. Hence, we are also interested in comparing the losses experienced by lenders when i's project fails.

First, we study the asymptotics of $w_{\cdot i}^*$. Recalling that $f_i(\cdot)$ is assumed to be exponential, having already shown that the intensity $\bar{F}_i(c_i^*)$ is of asymptotic order $\frac{1}{n\log(n)}$, (3.10) allows us to easily compute:

$$w_{\cdot i}^* = \frac{1}{\phi_i} - \Theta\left(\frac{1}{n\log(n)}\right).$$

Note that we must have $w_{\cdot i} < \phi_i^{-1}$ to ensure wealth remains positive, yet we can still pin down the rate at which the interbank investment approaches its upper bound.

[6] 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 Corollary 4.2 when the shock distribution is itself a power-law. The main result can be seen in Appendix B.

[7] We can obtain similar bounds using only Proposition 4.1, but these will not be tight. In particular, we would only show that $\bar{F}_i(c_i^*) = O\left((n\log(n))^{-\frac{\lambda_{i,L}}{\lambda_{i,U}}}\right)$, and $\bar{F}_i(c_i^*) = \Omega\left((n\log(n))^{-\frac{\lambda_{i,U}}{\lambda_{i,L}}}\right)$.

Next, we are interested in the term $\Gamma(\phi_i w^*_{\cdot i}; 1) = -\log(1 - \phi_i w^*_{\cdot i})$, which represents the relative loss of utility to a single lender when bank i's project experiences a failure. A straightforward computation using (3.3) gives

$$\Gamma(\phi_i w^*_{\cdot i}; 1) = \Theta\big(\log(n \log(n))\big) = \Theta\big(\log(n)\big),$$

since $\frac{\log(n\log(n))}{\log(n)} \underset{n\to\infty}{\to} 1$.

Putting this result together with the asymptotic rate of $\bar{F}_i(c^*_i)$, we see that

$$\bar{F}_i(c^*_i)(n-1)\Gamma(w^*_{\cdot i}; 1) = \Theta(1).$$

This is an interesting result, as it shows that the expected losses of utility due to a project's failure do not grow with the system size – in contrast, the decentralized setting exhibits $\bar{F}_i(\hat{c}_i)(n-1)\Gamma(\hat{w}_{\cdot i}; 1) = \Theta(n)$. Namely, the planner perfectly compensates for larger expected losses in utility through its reduction of a project's failure probability.

4.3 Price of Anarchy Asymptotics

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 that are lent a non-zero amount of capital in the planner's optimal allocation, i.e. $\mathcal{M}_n = \{i \in [1..n] : w^*_{\cdot i} > 0\}$. Banks in $\mathcal{M}_n$ form the 'core' of the financial network. If for some i we have $w^*_{\cdot i} = 0$, 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.

The 'price of anarchy' reflects how greedy decentralized behavior leads to lesser welfare in the system (Papadimitriou, 2001). In this model, we define it as

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

More precisely, the price of anarchy equals the relative loss in value between the centralized and decentralized settings. 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 $n \to \infty$, we have*

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

The proof is found in Appendix A.3, and uses all previous results from this Section.

It is particularly interesting that the price of anarchy does not grow with the system size n, or the remaining time horizon $(T-t)$. A more precise result can be obtained if banks are sufficiently homogeneous, where we can compute the limiting price of anarchy.

Corollary 4.4. *Assume 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 $|\mathcal{M}_n| \underset{n\to\infty}{\to} \infty$, then*

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

$$\frac{V}{n|\mathcal{M}_n|(T-t)} \underset{n\to\infty}{\to} \frac{\mu}{\phi}.$$

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

$$\text{PoA} \underset{n\to\infty}{\to} \frac{1}{1+\frac{\phi\theta\bar{F}(\hat{c})}{\mu}\left[\log\left(\frac{\phi\theta\bar{F}(\hat{c})}{\mu}\right)-1\right]}. \tag{4.1}$$

Corollary 4.4 verifies that the price of anarchy is of constant order 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 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 price of anarchy is increasing in $\frac{\phi\theta\bar{F}(\hat{c})}{\mu}$. Therefore, as the profitability of interbank loans in the decentralized setting is reduced, the limiting price of anarchy grows to infinity.

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

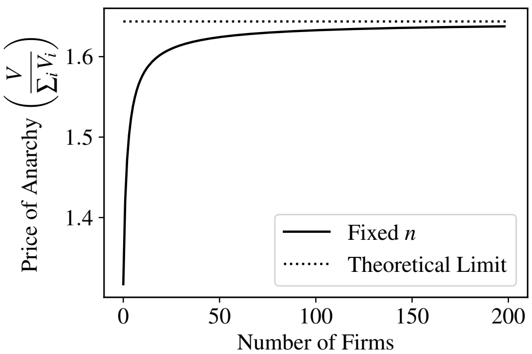

Figure 4.2: Simulating the Price of Anarchy for a system of identical firms as n grows.

Notation	Value	Description
r	0.01	Risk-free rate
μ	0.045	Excess drift
ϕ	0.4	Losses to lenders
η	0.5	Losses to borrower
$F(x)$	$1-e^{-\lambda x}$	CDF of shock size
λ	1	Parameter of $F(\cdot)$
θ	0.1	Shock arrival rate

Table 4.1: System parameters used in simulations for Figure 4.2. Code is available here.

4.4 Replicating the Centralized Allocation

Finally, we may be interested in studying how banks in the decentralized setting can be incentivized to replicate the planner's optimal allocation. To do so, we will allow the degree of each bank's investment into their own project, η_i, to vary. The rationale for this is twofold: first, η_i plays the fundamental role in decentralized banks' choice of how much cash to hold. If this value is sufficiently large, banks will increase their supply of liquidity and therefore reduce their projects' riskiness – which can lead them to meet the planner's optimal allocation. Second, we can imagine that lending banks are permitted to write a contract stipulating the borrowing 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 liquidity.

Let η_i^C (resp. η_i^D) denote the fraction of bank i's wealth lost upon project failure in the centralized (resp. decentralized) setting. We would like to choose η_i^D so that decentralized banks replicate the centralized optimum with 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:

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

First, notice that whenever $w^*_{\cdot i}(\eta_i^C) > 0$, the resulting value of η_i^D will grow exponentially in 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 risk. 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 lenders. 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 $w^*_{\cdot i}(0) > 0$. Namely, banks without counterparties hold no stake in their own projects.

5 Conclusions

In this paper we present a model by which banks in a financial system control both their own levels of risk, and their investment in each others' risky projects.

We compute the uniquely 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 interbank investments are closely related to a Sharpe-like ratio – which is controlled by borrowing banks. In particular, the optimal financial network exhibits a 'core-periphery' structure, wherein only a subset of banks serve as borrowers. Second, we formulate the optimization problem of a central 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. 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 counterparties 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 loss in welfare due to decentralized behavior grows with the size of both the financial system and its core. However, and perhaps counterintuitively, the *relative* loss of welfare, which we refer to as the price of anarchy, is of constant order. Finally, we show that it is possible, through regulation or contracting between banks, to replicate the planner's optimal interbank allocation. Banks who have borrowed the largest amount of capital will be subjected to the strictest requirements, and will therefore have the strongest incentive to reduce their project's riskiness. This highlights the danger of government bailouts, which can cause perverse incentives for individual banks.

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 a lending bank suffers losses of liquidity when their borrowers' project fails. 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 a 'lending linkage' to another bank is costless. It is natural to incorporate these costs into banks' optimization problems, for example, as the cost of performing due diligence on a borrower to assess their creditworthiness. Finally, the inclusion of intermediary costs or more sophisticated contracting mechanisms between banks presents a rich direction of future research.

Acknowledgments

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

References

Daron Acemoglu, Asuman E Ozdaglar, and Alireza Tahbaz-Salehi. Systemic Risk in Endogenous Financial Networks. *Columbia Business School Research Paper*, (15-17), 2015.

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.

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.

Sylvain Benoit, Jean-Edouard Colliard, Christophe Hurlin, and Christophe Pérignon. Where the Risks Lie: A Survey on Systemic Risk. *Review of Finance*, 21(1):109–152, 2017.

Sarah Bensalem, Nicolás Hernández Santibáñez, and Nabil Kazi-Tani. A continuous-time model of self-protection. *hal-02974961v2*, 2020.

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 Krahnen. Endogenous Banks' Networks, Cascades and Systemic Risk. *SAFE Working Paper*, 2014.

Stephen Boyd and Lieven Vandenberghe. *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.

Rodrigo Cifuentes, Gianluigi Ferrucci, and Hyun Song Shin. Liquidity Risk and Contagion. *Journal of the European Economic Association*, 3(2-3):556–566, 2005.

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.

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 Agathe Pernoud. Systemic Risk in Financial Networks: A Survey. *Annual Review of Economics*, 13:171–202, 2021.

Matthew O Jackson and Asher Wolinsky. A Strategic Model of Social and Economic Networks. *Journal of Economic Theory*, 71:44–74, 1996.

PA W Lewis and Gerald S Shedler. Simulation of Nonhomogeneous Poisson Processes by Thinning. *Naval Research Logistics Quarterly*, 26(3):403–413, 1979.

Deborah Lucas. Measuring the Cost of Bailouts. *Annual Review of Financial Economics*, 11(1): 85–108, 2019.

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 Possamaï. 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.

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

$$V_i(t,x) = \sup_{(c^i_\cdot, w^{i\cdot}_\cdot) \in \mathcal{A}^i_{t,T}} \mathbb{E}\Big[V_i\left(\tau, X^i_\tau\right) \Big| X^i_t = x \Big], \tag{A.1}$$

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

Next, we for each bank k we fix some admissible control $(c^k_\cdot, w^{k\cdot}_\cdot) \in \mathcal{A}^k_{t,T}$. 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:

$$\begin{aligned} V_i\left(\tau, X^i_\tau\right) - V_i\left(t, X^i_t\right) = & \int_t^\tau \left[\partial_t V_i(s, X^i_s) + \partial_x V_i\left(s, X^i_s\right) b_i(c^i_s, w^{i\cdot}_s) X^i_s\right] ds \\ & + \sum_{j=1}^n \int_t^\tau \left[V_i\left(s, X^i_s\right) - V_i\left(s, X^i_{s-}\right)\right] dN^j_s. \end{aligned} \tag{A.2}$$

where $b_i(c^i_t, w^{i\cdot}_t)$ is the coefficient on the dt term in (2.2).

Recall that the jump process N^j_t has instantaneous intensity $\theta_j \bar{F}_j(c^j_t)$. Therefore, the compensated process $M^j_t = N^j_t - \int_0^t \theta_j \bar{F}_j(c^j_s) ds$ is a martingale. Rewriting the integrals in (A.2) in terms of dM^j_t and taking expectation conditioned on $X^i_t = x$ (denoted $\mathbb{E}_{t,x}$) of both sides yields:

$$\begin{aligned} \mathbb{E}_{t,x}\left[V_i\left(\tau, X^i_\tau\right)\right] - V_i\left(t, X^i_t\right) = & \, \mathbb{E}_{t,x}\left[\int_t^\tau \mathcal{L}^{c^i_s, w^{i\cdot}_s} V_i(s, X^i_{s-}) ds\right] \\ & + \mathbb{E}_{t,x}\left[\int_t^\tau \left[V_i\left(s, X^i_{s-} - \eta_i X^i_{s-}\right) - V_i\left(s, X^i_{s-}\right)\right] dM^i_s\right] \\ & + \sum_{j \neq i} \mathbb{E}_{t,x}\left[\int_t^\tau \left[V_i\left(s, X^i_{s-} - \phi_j w^{ij}_s X^i_{s-}\right) - V_i\left(s, X^i_{s-}\right)\right] dM^i_s\right], \end{aligned} \tag{A.3}$$

where the generator $\mathcal{L}^{c_i, w_{i\cdot}}$ is defined to be

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

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:

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

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:

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

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

$$\sup_{(c_\cdot^i, w_\cdot^{i\cdot}) \in \mathcal{A}_{t,\tau}^i} \mathbb{E}_{t,x} \left[V_i \left(\tau, X_\tau^i \right) \right] = V_i \left(t, X_t^i \right).$$

Therefore, we arrive at:

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

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

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

Finally, applying the Dominated Convergence Theorem gives

$$0 = \sup_{(c_\cdot^i, w_\cdot^{i\cdot})} \mathcal{L}^{c_i, w_{i\cdot}} V_i(t, x),$$

which equals (3.2) after plugging in the definition of $\mathcal{L}^{c_i, w_{i\cdot}}$ from (A.4). □

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:

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

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

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

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 $c_i, w_{i\cdot}$. 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

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

Let $h(c_i, w_{i\cdot})$ 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

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

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 $\partial^2_{w_{ij}, w_{ij}} h < 0$. Additionally, $\partial^2_{c_i, c_i} h < 0$ because $\eta_j > 0$.

As a result, the Hessian matrix of the objective function is negative definite in the feasible region, i.e. $\nabla^2 h \prec 0$ 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 $y^* = (\hat{c}_i, w_{i\cdot}^*) \in \mathcal{A}_i$ satisfies

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

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

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

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:

$$\begin{aligned} \partial_t V_i(t,x) &= g_i'(t)\frac{x^{1-\gamma_i}}{1-\gamma_i} \\ \partial_x V_i(t,x) &= g_i(t)\frac{x^{1-\gamma_i}}{x} \\ V_i(t,(1-c)x) &= g_i(t)\frac{x^{1-\gamma_i}}{1-\gamma_i}(1-c)^{1-\gamma_i},\ \forall c < 1. \end{aligned}$$

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.

$$\begin{aligned} 0 = \frac{g_i'(t)}{1-\gamma_i} + g_i(t) \sup_{c_i, w_{i\cdot}} \Bigg\{ &(1-c_i)r + \sum_{j\neq i} w_{ij}\mu_j + \frac{\eta_i\mu_i}{\phi_i} + \theta_i \bar{F}_i(c_i)\frac{(1-\eta_i)^{1-\gamma_i}-1}{1-\gamma_i} \\ &+ \sum_{j\neq i} \theta_j \bar{F}_j(c_j)\frac{(1-\phi_j w_{ij})^{1-\gamma_i}-1}{1-\gamma_i}\Bigg\} \end{aligned}$$

$$g_i(T) = 1.$$

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:

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

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

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

Under Assumption 1, we will have both $\partial^2_{c_i, c_i} h < 0$ and $\partial^2_{w_{ij}, w_{ij}} h < 0$, 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 $\{c_s^i, w_s^{i\cdot}\}_{s \in [t,T]}$ 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:

$$\begin{aligned}
\psi(\tau^n, X^i_{\tau^n}) = \psi(t, X_t^i) &+ \int_t^{\tau^n} \mathcal{L}^{c_s^i, w_s^{i\cdot}} \psi(s, X_s^i) ds + \int_t^{\tau^n} \left[\psi\left(s, X^i_{s-} - \eta_i X^i_{s-}\right) - \psi\left(s, X^i_{s-}\right)\right] dM_s^i \\
&+ \sum_{j \neq i} \int_t^{\tau^n} \left[\psi\left(s, X^i_{s-} - \phi_j w_s^{ij} X^i_{s-}\right) - \psi\left(s, X^i_{s-}\right)\right] dM_s^j.
\end{aligned}$$

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

$$\begin{aligned}
\mathbb{E}_{t,x}\left[\psi(\tau^n, X^i_{\tau^n})\right] = \psi(t,x) &+ \mathbb{E}_{t,x}\left[\int_t^{\tau^n} \mathcal{L}^{c_s^i, w_s^{i\cdot}} \psi(s, X_s^i) ds\right] \\
&+ \mathbb{E}_{t,x}\left[\int_t^{\tau^n} \left[\psi\left(s, X^i_{s-} - \eta_i X^i_{s-}\right) - \psi\left(s, X^i_{s-}\right)\right] dM_s^i\right] \\
&+ \sum_{j \neq i} \mathbb{E}_{t,x}\left[\int_t^{\tau^n} \left[\psi\left(s, X^i_{s-} - \phi_j w_s^{ij} X^i_{s-}\right) - \psi\left(s, X^i_{s-}\right)\right] dM_s^j\right].
\end{aligned}$$

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

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

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

$$\mathbb{E}_{t,x}\left[U_i(X_T^i)\right] = \psi(t,x) + \mathbb{E}_{t,x}\left[\int_t^{T} \mathcal{L}^{c_s^i, w_s^{i\cdot}} \psi(s, X_s^i) ds\right]. \tag{A.11}$$

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

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

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

$$\psi(t,x) \leq \sup_{\{c_s^i, w_s^{i\cdot}\}_{s\in[t,T]}} \mathbb{E}_{t,x}\left[U_i(X_T^i)\right] = V_i(t,x). \tag{A.12}$$

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

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

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

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

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. □

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 $V(t, X_t^1, \ldots, X_t^n)$ yields more terms, but remains simple as the jump processes are mutually independent. Namely, the generator is given by

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

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

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

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

and conclude as in the previous result. □

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 $\left(f_i(x) = \frac{(\alpha_i^{-1}-1) x_0^{\alpha_i^{-1}-1}}{(x+x_0)^{\alpha_i^{-1}}}, \text{ for any } x_0 > 0 \text{ and } \alpha_i < 1 \right)$. 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 $w_{\cdot i}^* = 0$ – as shown in (3.10). Of particular interest, the optimal decentralized and centralized allocations for c_i and $w_{\cdot i}$ will coincide whenever either $c_i^* = 0$ or $w_{\cdot i}^* = 0$ in the planner's optimum.

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

$$\begin{aligned} 0 = \partial_t V + \sup_{c_\cdot, w_{\cdot\cdot}} \Bigg\{ \sum_{i=1}^n \Bigg(& \left[(1-c_i) r + \sum_{j\neq i} w_{ij}\mu_j + \frac{\eta_i \mu_i}{\phi_i} \right] x_i \partial_{x_i} V \\ & + \theta_i \bar{F}_i(c_i) \Big[V(t, x_1(1-\phi_i w_{1i}), .., x_i(1-\eta_i), .., x_n(1-\phi_i w_{ni})) - V \Big] \Bigg) \Bigg\} \qquad \text{(A.14)} \\ V(T, x_1, ..., x_n) = \sum_{i=1}^n U_i(x_i). & \end{aligned}$$

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: $V(t, x_1, .., x_n) = g(t) + \sum_i \log x_i$. Substituting into (A.14), we obtain:

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

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 lend the same fraction of their wealth to bank i.[8] Let this fraction be denoted by $w_{\cdot i}$. This allows us to further simplify (A.15) and obtain

$$0 = g'(t) + \sum_{i=1}^{n} \frac{\eta_i\mu_i}{\phi_i} + \sup_{c_\cdot, w_{\cdot\cdot}} \sum_{i=1}^{n} (1-c_i)\, r + (n-1)w_{\cdot i}\mu_i - \theta_i \bar{F}_i(c_i)\Big[\Gamma(\eta_i; 1) + (n-1)\Gamma(\phi_i w_{\cdot i}; 1)\Big].$$

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

$$\sum_{i=1}^{n} \sup_{(c_i, w_{\cdot i})\in\mathcal{A}_i} h_i(c_i, w_{\cdot i}), \tag{A.16}$$

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

Reduction to Univariate Optimization: We first maximize over $w_{\cdot i}$ and then c_i given the optimal $w_{\cdot i}$. Given a value of c_i, we seek to find the optimal value of $w_{\cdot i}$. We can compute

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

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:

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

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

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

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

[8] 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.

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

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

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

As $c_i \to \infty$, we will have $\bar{F}_i(c_i) \to 0$. 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 $x \log x \underset{x \to 0}{\to} 0$, we will have $\lim_{c_i \to \infty} h_i^*(c_i) = -\infty$.

This limit is sufficient for existence of an optimal solution to (A.19). Fix some $K < 0$. Since we have shown $h_i^*(c_i) \underset{c_i \to \infty}{\to} -\infty$, 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 compliment. 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

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

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:

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

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

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

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:

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

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

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

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 $h_i^*(\cdot)$:

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

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

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

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

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

Applying Grönwall's inequality, we see that

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

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

Rewriting (A.22), we obtain:

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

For i satisfying

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

in the assumption (3.9), we see that $\frac{d^2h_i^*}{dc_i^2}(\tilde{c}_i) < 0$. 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

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

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 $\frac{d^2h_i^*}{dc_i^2}(c_i) < 0$ 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 $(\tilde{c}_i, \infty)$. At such a point c_i^*, we must have $\frac{dh_i^*}{dc_i}(c_i^*) = 0$, and also $\frac{d^2h_i^*}{dc_i^2}(c_i^*) \le 0$.[9] By the same conclusion using Grönwall's inequality, we must have $\frac{d^2h_i^*}{dc_i^2}(c_i) < 0$, 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. □

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 $\{c_s^{\cdot}, w_s^{\cdot\cdot}\}_{s\in[t,T]}$. We then apply Itô's formula, which only differs in yielding a few more terms.

[9]These are the two necessary conditions for optimality of c_i^* when it lies in the interior of the feasible region.

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 $\left\{X_s^1\right\}_{s\geq 0} \dots \{X_s^n\}_{s\geq 0}$, is bounded at time τ^n. Therefore, we choose

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

and conclude identically. $\square$

A.3 Differences in Optima

Proof of Proposition 4.1. The main idea in this proof is to first establish crude bounds of:

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

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

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

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

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

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

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

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

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

where $K = \max_i \left\{\hat{c}_i + \frac{\theta_i}{r}\bar{F}_i(\hat{c}_i)\big[1 + \Gamma(\eta_i;1) + \Gamma(\phi_i \hat{w}_{\cdot i};1)\big]\right\}$ does not depend explicitly on n. Since $w_{\cdot i}^* \geq 0$, 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

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

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:

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

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

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

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 $c_i^* \leq Kn^2$, which, when plugged in, yields:

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

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

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

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:

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

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 $\lim_{n\to\infty} \frac{c_i^*}{\log(n)} \leq \lambda_{i,U}$.

(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

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

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

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

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

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

This is precisely the lower bound in the first part of Proposition 4.1. 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:

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

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

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

Let $C_L = \frac{\kappa_{i,U}\lambda_{i,U}}{\kappa_{i,L}\lambda_{i,L}} \left(\frac{r}{\theta_i} \right)^{\frac{\lambda_{i,L}}{\lambda_{i,U}} - \frac{\lambda_{i,U}}{\lambda_{i,L}}} \frac{\left(\kappa_{i,U}\Gamma(\eta_i;1)\right)^{\frac{\lambda_{i,U}}{\lambda_{i,L}}}}{\left(\kappa_{i,L}\Gamma(\phi_i \hat{w}_{\cdot i};1)\right)^{\frac{\lambda_{i,L}}{\lambda_{i,U}}}}$. 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 $-\log\left(\frac{\phi_i\theta_i}{\mu_i}\bar{F}_i(\hat{c}_i)\right) = \Gamma(\phi_i\hat{w}_{\cdot i};1) \geq 0$, and hence this term can be dropped. Simplifying, we arrive at the desired bound of:

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

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

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

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

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

where V and V_i are evaluated at $(t, x_1, ..., x_n)$ and the difference becomes independent of wealths because of logarithmic utility. Notice that any of the terms in the sum will equal zero if $w_{\cdot i}^* = 0$ (in which case we also must also have $\hat{w}_{\cdot i} = 0$, and hence $\hat{c}_i = c_i^*$). If not, then using the results from Section 4 we see that

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

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

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

In Proposition 3.2, it is easy to see that $V_i = (T-t)\Theta\left(|\mathcal{M}_n|\right)$, 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 $\hat{c}_\cdot$ indicates the decentralized optimum, and $c_\cdot^*$ indicates the centralized optimum (likewise for $w_{\cdot\cdot}$).

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 $n \to \infty$. Moreover, by assumption that all banks in $\mathcal{M}_n$ are homogeneous, we will have $\hat{w}_{ij} = \hat{w}_{ik}$ for any $j, k \in M_n$. This yields:

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

where $\hat{c}$ denotes the optimal liquidity supply held by 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:

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

and the desired limit follows.[10]

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:

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

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 $(n-1)\bar{F}_i(c_i^*)\Gamma(\phi_i w_{\cdot i}^*;1)$ 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 $w^* \to \phi^{-1}$ as $n \to \infty$, 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. □

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 Proposition 4.1:

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:***

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

[10] We note that this expression for J_i^* is only correct when i is not in $\mathcal{M}_n$, otherwise we would have a factor of $|\mathcal{M}_n| - 1$ in front of the term in brackets. However, in the limit this difference will vanish.

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

(ii) ***Lower Bound:***

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

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

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 Corollary 4.2.

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

$$c_i^* = \Theta\left(\left[(n-1)\log(n)\right]^{\alpha_i}\right).$$

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 Proposition 4.1. Recall that we have shown that

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

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

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

and

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

We can then follow the proof of Proposition 4.1 identically, but using these bounds instead. □